

THIS REPORT HAS BEEN DELIMITED
AND CLEARED FOR PUBLIC RELEASE
UNDER DOD DIRECTIVE 5200.20 AND
NO RESTRICTIONS ARE IMPOSED UPON
ITS USE AND DISCLOSURE.

DISTRIBUTION STATEMENT A

APPROVED FOR PUBLIC RELEASE,
DISTRIBUTION UNLIMITED.

BEST COPY
AVAILABLE

Armed Services Technical Information Agency

Because of our limited supply, you are requested to return this copy WHEN IT HAS SERVED YOUR PURPOSE so that it may be made available to other requesters. Your cooperation will be appreciated.

AD

40840

NOTICE: WHEN GOVERNMENT OR OTHER DRAWINGS, SPECIFICATIONS OR OTHER DATA ARE USED FOR ANY PURPOSE OTHER THAN IN CONNECTION WITH A DEFINITELY RELATED GOVERNMENT PROCUREMENT OPERATION, THE U. S. GOVERNMENT THEREBY INCURS NO RESPONSIBILITY, NOR ANY OBLIGATION WHATSOEVER; AND THE FACT THAT THE GOVERNMENT MAY HAVE FORMULATED, FURNISHED, OR IN ANY WAY SUPPLIED THE SAID DRAWINGS, SPECIFICATIONS, OR OTHER DATA IS NOT TO BE REGARDED BY ANY PERSON OR CORPORATION, OR CONVEYING ANY RIGHTS OR PERMISSION TO MANUFACTURE, REPRODUCE, OR SELL ANY PATENTED INVENTION THAT MAY IN ANY WAY BE RELATED THERETO.

Reproduced by
DOCUMENT SERVICE CENTER
KNOTT BUILDING, DAYTON, 2, OHIO

UNCLASSIFIED

TO NO. 4222
STIA FILE COPY

TECHNICAL REPORT NO. 2

A Survey of
Bernoullian Utilities & Applications

by

Ernest W. Adams

BUREAU OF APPLIED SOCIAL RESEARCH
COLUMBIA UNIVERSITY

Behavioral Models Project
(NR 012-115)

TECHNICAL REPORT NO. 9

A Survey of Bernoullian Utilities and Applications

by

Ernest W. Adams

June 1954

CU-11-54-NONR-266(21)-BASR
Bureau of Applied Social Research
New York 27, N.Y.

Table of Contents

Page

A Survey of Bernoullian Utilities and Applications

1. General Introduction to Utility Theory.	1
1.1 Informal Description of Utility.	1
1.2 Formalization of Utility Theory.	14
1.3 Questions of Interpretation and Confirmation	26
1.3.1 Interpretation and Confirmation in the Descriptive Theory of Utility	30
1.3.2 The Definition of Rationality.	35
1.3.3 The Definition of Good	37
2. Bernoullian Utilities	39
2.1 Introduction: the Problem of Rationality	39
2.2 Bernoullian Utility Functions.	42
2.2.1 The Defining Conditions.	42
2.2.2 The Primitive Notions and Their Interpretations	45
2.3 Axioms	47
2.4 A Different Approach	57
2.5 The Mapping Theorem; the Construction of the Bernoullian Utility Function; the Uniqueness of the Function. . .	61
3. The Decision Problem.	67
3.1 Description of the Decision Problem and its Relation- ship to Utility Theory.	67
3.2 Formalization of the Decision Problem.	71
3.3 Mixed Strategies	77

	<u>Page</u>
3.4 Strategies in Statistical Games.	79
3.4.1 Payoff Function and Strategies for Single- Experiment Games	83
3.4.2 Sequential Games	89
3.5 Decision Principles.	96
3.5.1 Two-Person Games Between Two Rational Players: the Minimax Principle.	98
3.5.2 Strategies in n-Person Games, All Players Rational	108
3.5.3 Statistical Games.	110
4. Descriptive Applications of Bernoullian Utility Theory.	114
4.1 Hypotheses Explaining Gambling, and Insurance-Buying . .	116
4.2 The Mosteller-Weges Experiment	126
4.3 Problems of Interpretation and Confirmation: the Future of Bernoullian Utility Theory.	134
Bibliography.	136

A SURVEY OF BERNOULLIAN UTILITIES AND APPLICATIONS

1. General Introduction to Utility Theory

1.1 Informal Description of Utility

The concept of utility has had a career in economic theory dating at least from Adam Smith,¹ during the course of which it has undergone many important modifications of meaning. Before entering into the details of the specific forms that the utility concept has taken, we shall try to indicate the common core of meaning, and the types of problems in which this concept has been used.

Let A be an individual who is at a given time presented with the necessity of choosing among a set of alternatives, E_1, E_2, \dots, E_n . To take a specific example, suppose A is Mr. Jones who is at the market and is considering which of the following three items to buy: a steak, four bottles of milk, or a bottle of wine, which are E_1, E_2 , and E_3 respectively. To a certain extent Mr. Jones' choice will be determined by the prices of the items, and the amount of money he has, but to a certain extent also his choice will be determined by the value of these commodities to him. If the prices of the three alternative items, E_1, E_2 , and E_3 are the same, then we may very well expect that his choice will depend solely on his valuation of the commodities. Another term frequently used for this subjective valuation of the different alternatives is utility. In the example above, Mr. Jones' choice depends on the prices, and on the utility of the items for him, and in the case in which the prices are all equal, he will choose that item with the greatest utility.

1. See Stigler, G. S. [29]

The example given above is typical of the way in which the notion of utility enters into a large area of problems. These problems all involve, in one way or another, an element of choice, made by one or more individuals, among a set of several alternatives. It is then usually assumed that each individual possesses a utility scale by which he ranks the alternatives according to greater or lesser utility, and that the actual choice made by him depends in some fixed way upon the utility rankings of the alternatives presented.

When we examine our original example more closely, we notice some points which bring out some of the major differences between various concepts of utility. We stated that Mr. Jones had three alternatives: to buy a steak, four bottles of milk, or a bottle of wine. But in most cases, a consumer is not faced with that sort of choice. He can usually, within the limits of his budget, buy all of the items, or any combinations of them which suits his fancy, or none at all. Hence in order to take full account of Mr. Jones' preferences, we must include in the set of alternatives, E_1, \dots, E_n , all the available courses of action he can possibly take on this occasion. In our example, then, we must include not only the utilities of steak, milk, and wine, but the utilities of steak and milk, steak and wine, etc. In the past it was frequently assumed by economists that in order to obtain the utility of a combination of two items, such as steak and wine, it was sufficient simply to add the separate utilities of the items arithmetically. This particular assumption implies some rather special assumptions about the nature of an individual's utility scales, and these were increasingly criticized until the assumption of additive utilities was finally abandoned.

Two simple examples should be enough to convince the reader that at least in certain cases this hypothesis is absurd. Let it be required to find the utility of a combination of a phonograph and a collection of records. Clearly this cannot be the sum of the utilities of the phonograph and records separately, for each without the other has no value. In this case we may say that the two items complement each other. In other cases pairs of items may compete with each other, as for example, wrist-watches and pocket-watches; that is, the utility of the two together may be less than the sum of their separate utilities. Much controversy in the past centered in attempts to define independent sets of commodities, that is, sets of commodities for which the utility of a combination is the sum of the utilities of the elements of the combination.

A second example contradicting the hypothesis arises when the combination consists of a number of units of the same item. Under the hypothesis of additive utilities, the utility of n loaves of bread must be $u + u + \dots + u$ (n times) $= nu$ where u is the utility of one loaf of bread. However, most people would deny that a thousand loaves of bread are a thousand times as valuable to them as one loaf. This example is, of course, a special case of competing commodities. Here the loaves of bread, like the wrist-watch and the pocket-watch, compete with each other in the sense that one, or at most, a few loaves of bread satisfy the customer's needs, and the remainder have very little additional utility. By applying similar reasoning to the demands for all commodities, economists were led to the principle of diminishing marginal utility, which has played a prominent part in the classical analysis of consumer behavior. The principle of diminishing marginal

utility states that the utility increase with each additional unit of a given item becomes smaller the larger the total amount of that item already present. While this formulation of the principle is open to objections which will be brought out later on, it can be reformulated to meet them and yield empirical implications in a field which is otherwise rather barren of them.

Going back to Jones again, we specified that his choice depends partly on his utility scale, but partly also on the prices and on how much money he has. As long as we do not specify precisely how his choice depends on the utility scale, we have a theory without predictive value. In the classical theory of consumer behavior,¹ the consumer was usually assumed to spend his money in such a way that the set of items purchased had the greatest utility of all the sets of goods which could have been purchased within the customer's income at the given prices. Thus, in the classical theory three variables were involved: prices, utilities, and income, and it was the task of classical utility theorists to discover how changes in any of these would affect the consumer's buying pattern. It is possible, and it has been done by many modern theorists, to take a different tack, reducing the set of relevant variables from three to two. Again the consumer, or any agent confronted with a choice, chooses among a set of possible alternative courses of action, which may, as before, include buying steak, milk, or wine. However, in assessing the utility of a particular alternative, say buying a steak, he considers the utility of the entire act of purchase. This includes not only receiving the steak, but paying over the price demanded. The difference between the two types of analysis lies in the fact that in the first case, the valuation

1. See Samuelson, P. A. [25] pp. 90-121.

5

of the particular alternative, depends only on the individual's liking for steak, whereas in the second case his valuation must include both his expected satisfaction from the steak and his valuation of the money to pay for it. Therefore the price does not enter directly in the second case as a variable which determines choice.

Obviously this suppression of the variables of price and income does not simplify the problem of explaining and predicting consumer's behavior, since these variables still affect it, only now by way of affecting the utility scale itself, which in the other analysis had been considered independent of price and income.

Our formulation of the problem of choice has now reached substantially the standard modern form: every individual has a utility scale by which he ranks all things, and when presented with a choice among a set of possible courses of action, he chooses that alternative which is highest in his utility scale.

So far our discussion has shown us that one of the major differences between various utility concepts lies in the type of entity which is taken to be evaluated in the individual's utility scale. In our original example, certain individual items of consumption were ranged in order on Mr. Jones' utility scale. However, it was found necessary to include not only individual items, but also all the possible combinations of the basic items. Later it was suggested that the alternatives ranked should be not simply the possible bundles of items to be bought, but the total value of the transactions of buying including the value of the money payment as well. These three kinds of alternatives do not exhaust the possibilities. If a theory of choice is to

encompass choices made in all situations, it must include among the entities ranked all the kinds of alternatives which may be encountered by an individual in the process of making a decision. So far we have mentioned only the individual making decisions in his capacity as consumer. But of course individuals make decisions in other than buying situations, and in situations where the exchange of money is only an incidental feature, such as whether to go to a movie or stay home and work. In some cases, it seems reasonable to consider as the alternatives to be ranked not the particular acts which would be carried out as the consequence of a decision but the future history of the individual which he expects to be consequent on his decision. For example, in evaluating the utility of a bottle of wine, the individual would consider not only his liking for wine, but all the consequences he considers likely to follow from its purchase. Of course, this shift from considering not only the immediate consequences of a decision but all the expected consequences through time is only a change of terminology, since most people tacitly include these in determining the utility of an alternative. Some such considerations must be involved in the calculation of the utility of losing a dollar in paying for any item, since for most people the value of a dollar lies only in what it can be used for. Taking the set of alternatives to be possible histories emphasizes the fact that the utility of any particular decision depends not only on the act of buying or the immediate satisfactions of the purchase, but on all the consequences expected to accrue from it.

From the consideration of histories as relevant alternatives we are led to consider still another kind of prospect: alternatives involving uncertainties, probabilities, or risks. If, in evaluating the utility of say, buying a car, a person must take into account all the consequences of

this purchase, he must consider a number of possibilities - such as being involved in various accidents - about which he can make no certain predictions. Particularly clear-cut examples of risk situations arise in gambling, in which all the probabilities may be known exactly. Suppose that a man is trying to decide whether or not to buy a lottery ticket costing one dollar with probability α of winning n dollars and probability $1-\alpha$ of losing the dollar he pays for the ticket. He must then compare the utilities of two different futures in his effort to decide whether or not to buy the ticket. The future involved in not buying the ticket is certain, at least with respect to the outcome of the lottery. The future involved in buying the ticket is an uncertain combination of two other certain futures: the future consequent on losing the dollar and the one consequent on winning the n dollars.

The inclusion of uncertain futures as well as certain ones among the set of prospects or alternatives to which the individual assigns utilities would be of little interest were it not for the fact that using some very plausible assumptions, a very simple relation may be established between the utilities of sure future prospects and the utilities of uncertain combinations of these prospects. From these assumptions, it follows that the utility of a combination of two prospects, E_1 and E_2 with utilities u_1 and u_2 respectively, combined into a single uncertain prospect of E_1 with probability α and E_2 with probability $1-\alpha$, is simply the expected value of the utilities of these prospects: $\alpha u_1 + (1-\alpha) u_2$. Utility scales which have the property that the utility of a probability distribution of sure prospects is the expected value of the utilities of those prospects are called Bernoullian utilities after Daniel Bernoulli who was the first to make such an assumption.

Throughout the foregoing discussion, we have assumed that the agent confronted with a decision is a single person. For psychological reasons,

perhaps, it seems most reasonable to apply utility theories to individuals, since we feel that we have some insight into the process of decision making by them. There is, however, no logical reason why a utility type of analysis cannot be extended to other types of agents confronted with the necessity of making decisions, such as families, business enterprises, and governments. In the theory of consumer behavior, the basic consumer unit for many purposes is taken to be not the single person but the household he represents. In this case, it is appropriate to talk of the utility scale of the household, and of buying to maximize the utility of the household within the limitations of the household's income. The fact that we assign utility scales to organizations composed of many people who will in general have utility scales of their own brings up another problem which we only mention here. This is the question of how the utility scale of an aggregation is related to the utility scales of the individuals composing it.

Let us recapitulate our utility theory and some of its related problems. The theory involves an agent, A , confronted with making a decision among a certain set of possible alternatives, E_1, \dots, E_n , these being variously interpreted as actions or as the outcomes of an action. A ranks the alternatives according to a utility scale, and selects that with the highest utility. We have seen that the individual, A , may stand for different sorts of entities, both human and institutional, and that the set of alternatives, E_1, \dots, E_n , may also be differently interpreted in different types of utility theory. One further question, not so far raised, is that which asks what sort of thing the utility scale is.

We can go some distance in our answer to the last question before arriving at the limits of controversy. Again, let A be an agent confronted

with a decision among alternatives E_1, E_2, \dots, E_n . We assume that A ranks the alternatives in some manner by preference; that is, for any two alternatives, E_i and E_j , A is able to say either that he prefers E_i to E_j or E_j to E_i , or that he is indifferent between them. One condition placed on the utility scale must certainly be that it reflect A's preference pattern; that is, if one alternative is preferred to another, the first must have greater utility than the second, and if the two alternatives are equally preferred their utilities must be equal. We can formalize this condition by saying that a utility scale is a function, which we designate u , which is defined for the set of alternatives E_1, \dots, E_n , such that for all $i = 1, \dots, n$, $u(E_i)$ is a real number, and which must satisfy the following condition:

(1) for all $i, j = 1, 2, \dots, n$,

$$u(E_i) > u(E_j)$$

if and only if E_i is preferred to E_j . This condition simply states in a formal way that the utility function, u , reflects the individual's preference scale.

There is no controversy in the characterization of the utility function to this point, and it is worth pointing out that even such an apparently trivial condition as (1) has some empirical consequences. The most important consequence is that the individual's preference ordering of the alternatives must be transitive, i.e., if E_i is preferred to E_j and E_j to E_k , then E_i must be preferred to E_k . If this were not the case, and there existed some "preference circles" such as E_i preferred to E_j , E_j to E_k , and E_k to E_i , then no function u could exist satisfying condition (1). The requirement of transitivity is often referred to as the requirement of consistency.

However, requirement (1) is rather weak, since if u is any function satisfying (1), any other function v , which satisfies the condition that

for all x and y ,

(2) $v(x) > v(y)$ if and only if $u(x) > u(y)$

also satisfies condition (1). Thus, for example, the functions $2u$, $u^3 + 3$, and e^u also satisfy condition (1). Any two functions satisfying condition (2) are said to be monotonically related, or one is called a monotonic transformation of the other. Condition (1) is satisfied by every monotonic transformation of u if it is satisfied by u , and we say that condition (1) defines a utility function uniquely only up to a monotonic transformation. In general, if a condition is given which defines a function, u , only up to a monotonic transformation, the only thing of significance about the values of u are the relative magnitudes of $u(x)$ and $u(y)$ for any two arguments, x and y , for which the function is defined. The absolute magnitude, $u(x)$, or the numerical value of the difference $u(x) - u(y)$ is generally without significance, since we can always replace u by another monotonically related function v and have $v(x)$ and $v(x) - v(y)$ assume arbitrary values (as long as $v(x) - v(y)$ has the same sign as $u(x) - u(y)$). Throughout the history of economics, other conditions have been placed on the utility function, but condition (1) is the only one on which there has been general agreement. Those economists who came to believe that (1) is the only meaningful condition to be placed on the utility function were often led to the conclusion that it would be better to discard the utility function entirely, and work directly with the individual's preference pattern, since the utility function tells us no more than the preference function, and has the psychological disadvantage of appearing to contain more significance than it actually has. This position known as Ordinalism because it holds that the only significance of the utility function is the ordering it assigns to the alternatives according to their utility values. Contrasted

with the "ordinalist school" are various "cardinalist schools" which by placing additional restrictions on the utility function, are able to define utility functions with more significance than the ordering they assign to the alternatives.

Because the utility functions defined by different sets of conditions often resemble each other in mathematical respects, it has sometimes been assumed that the functions defined by two different sets of conditions are the same. Logically there is no reason why this should be true, and if it is true, it stands in need of a rigorous justification which is not usually given. This last remark will be amplified below.

The early economists' assumption that the utility of a combination of items is equal to the sum of the utilities of the items was mentioned earlier. We can express this assumption in terms of the utility function, u , as follows: let E_1 and E_2 be two distinct consumption items, and let $E_1 \& E_2$ be the item which consists of E_1 and E_2 together. Then

$$(3) \quad u(E_1 \& E_2) = u(E_1) + u(E_2).$$

Condition (3) on the utility function is clearly much stricter than condition (1) in the sense that many functions which satisfy (1) do not satisfy (3).

The general problem of determining the set of functions satisfying (3) is not completely solved; however, with a few additional plausible assumptions (including condition (1)) it can be shown that the functions satisfying these assumptions are unique up to multiplication by a positive constant. That is, if u satisfies these conditions, then the only other functions v which also satisfy these conditions must satisfy the equation

$$(4) \quad v = ku$$

for some positive number k .

A set of conditions which restrict the utility function so narrowly as those satisfying equation (4) above are said to yield a cardinal measure of utility. A parallel example of a cardinal measure is the case of measurement of physical mass, in which the actual value of the mass measurement for any particular body is determined once a unit of measurement is fixed. Similarly, under the conditions mentioned above, the measurement of utility is uniquely determined once a unit of measurement (often called a utile) is fixed on. Most of the early economists assumed that utilities, like most of the physical measurements known to the science of their time, were cardinally measurable. From this assumption it was easy for them to take a still further step and assume that the utility measures of different individuals were comparable. That is, it was assumed that it is meaningful to speak, for example, of a given alternative as having twice as much utility for individual A as for individual B. Under a utilitarian system of ethics, in which ethical good is based on individual utilities, the interpersonal comparability of utility scales would make it possible to combine the utilities of individuals so as to obtain a total social utility which could be made the basis of social policy recommendations.

This last application of utility theory is properly a part of welfare economics, which is that part of economics which takes for its task the recommendation of social policy in the economic sphere. Because of the utilitarian bent of prevailing social philosophy in England and the United States, theories of welfare economics in these countries have often been based on an underlying utility theory.

On page 11 above, we alluded to certain difficulties in defining a social welfare function based on the cardinal utilities of individuals.

In general, cardinal measures may be defined in several ways, depending on what conditions the utility function is expected to satisfy. As we have seen, condition (3) with a few additional assumptions defines a cardinal measure of utility. The general theory of utilities with risks leads to another cardinal measure. Which is to be taken to be in some sense as the measure of the individual's good? One should note on reading over the conditions placed on the utility function that these do not necessarily guarantee that the function satisfying the conditions yields a true measure of the amount of satisfaction which the individual expects to gain from the ranked alternatives. The possibility that a given cardinal utility measure may not be a measure of the individual's good has been grounds for criticism of many proposed cardinal utility measures.¹ Of course, this criticism would be pointless if the cardinal utility in question were intended only to predict the behavior of individuals, or to predict general consumer trends; often, however, the main reason for constructing cardinal utility scales has been in order to use them as a basis for policy recommendations.

Other schools of welfare economics attempt to build social policies on ordinal utilities alone. A rather simple use of these in this connection is in voting, in which, in its simplest form, each individual indicates which of two alternatives he prefers. Other more complicated schemes have been considered, which will be discussed in section 3. The question of the legitimacy of comparing ordinal utility scales of individuals may arise here just as it arises with respect to cardinal utilities.

We have now briefly outlined the main areas of application of utility theory. One area involves the problems of explaining and predicting individual

1. See e.g., Vickroy [30]

behavior in choice situations on the assumption that this behavior is in accordance with a utility scale. In this area is included a variety of theories which are roughly subdivided as ordinal and cardinal according to the conditions they place on the utility functions. The second major area includes application to problems involving the aggregation or comparison of individual utilities. In this area too, the theories may be based on either ordinal or cardinal utilities. A second way of subdividing this area is according to whether the theory is normative or descriptive. Welfare economics, as a theory providing regulative (i.e., normative) principles, falls in the first of these categories. Until recently welfare economics was the only discipline making use of utility in this area. There are now, however, some theories which fall into the descriptive category of this second area. These theories attempt to describe the way in which interaction influences utilities, and in what way utilities of individuals must be aggregated to form a group utility meeting certain specifications. This report deals with utilities of individuals, and thus is not concerned with aggregation problems and welfare economics.

1.2 Formalization of Utility Theory

We have now reached a point in our discussion at which it may be profitable to introduce a few formal mathematical notations for some of the important concepts of utility theory. The reader will be assumed to be familiar with such elementary mathematical notions as those of a set, class membership relation, function, real number, and the standard mathematical notations for these.

We shall be concerned with the utilities and preferences of one individual. X will denote the set of alternatives which are ranked by the

individual. Previously we have denoted the alternatives by E_1, \dots, E_n ; however, we do not wish to limit the alternatives to a finite or even countably infinite number. In case we want to include probability distributions of sure alternatives, there must be a continuum of alternatives; therefore K is an arbitrary non-empty set.

The individual's preference-or-indifference relation is denoted by " \succeq ". Thus if x and y are elements of K , then

$$x \succeq y$$

means that the individual either prefers x to y or is indifferent between x and y . Henceforth $x \succeq y$ will be read as 'x is preferred or indifferent to y.' We have taken the relation of preference-or-indifference as basic because both the relation of preference and the relation of indifference are definable in terms of this one relation, R . We define the indifference relation for the individual, denoted " \sim " as follows: for all x and y in K ,

$$(6) \quad x \sim y \text{ = df } x \succeq y \text{ and } y \succeq x.$$

The preference relation, " $>$ " is defined in terms of " \succeq " and " \sim " for all x and y in K ,

$$(7) \quad x > y \text{ = df } x \succeq y \text{ and not } x \sim y^{(1)}$$

Finally, " u " denotes the individual's utility function. Then, for all $x \in K$, $u(x)$ is a real number.

We shall take K , \succeq , and u as basic notions in an axiomatic treatment of utility.

(1) This is read: "x is preferred to y" and means the individual either prefers x to y or is indifferent between them, but he is not indifferent between them; or in the simplified reading, x is preferred-or-indifferent to y, but x is not indifferent to y.

We are now in a position to restate formally some of the fundamentals of utility theory, and some of the defining features of its variants. In the theory of consumer behavior with which we began our discussion of utility theory, we are concerned with a single individual. It is true, that in order to be useful to economics, the consumer theory must be applicable to a large number of people, but in this case only as a statistical aggregate of the behavior of independent agents. Therefore, as a first approximation, the problem of predicting the behavior of an individual is solved for individuals who are assumed to act independently, and from these the consumption patterns of the community are derived, using suitable assumptions about uniformity of individual tastes.

Differences between approaches to consumer behavior appear in the different types of entities taken to be members of K . Each interpretation for the class K leads to a different type of utility theory. Below are listed 5 different interpretations of K , and some indication of the types of theories with which they are associated.

(a) K is the set of commodity bundles. This leads to a theory based on independent variables of price and income which together determine the set of possible alternatives from which consumers may choose. This is the usual interpretation in the classical theory of consumer behavior.

(b) K consists of commodity bundles together with their prices. No existing theories are based directly on this interpretation, but this is an intervening stage between interpretations (a) and (c). A consumer theory based on interpretation (b) would be much like the classical theory: in fact, classical theory can be interpreted in this way, where the utility of money is included with the utility of any commodity bundle to determine its total utility.¹ The conceptual difference between (a) and (b) is that in (b) we

1. See for example, Samuelson, [25] p. 99

consider the total utility of any alternative act, which, if the act is buying, includes payment of the price.

(c) K consists of total histories expected to be consequent on any given decision. The only difference between (b) and (c) is that in (c) attention is focused on the consequences of the decision, and in (b) we appear to be concerned with the decision itself. We would expect that the utility arrived at in each case would be the same, whether we were to consider the actual decision, or the consequences of the decision.

(d) K contains, in addition to sure alternatives like those of (b), probability distributions over these alternatives. With suitable axioms, this interpretation leads to the theory of Bernoullian utilities mentioned on page 7. This theory is then used to handle problems of individual reactions to risk, such as in games, gambling, and purchase of insurance.

(e) K contains histories as in (c), and probability distributions over histories. Like the difference between interpretations (c) and (b), interpretations (c) and (d) do not differ as much in content as in emphasis.

In parts 2 - 4 of this report we shall be concerned principally with theories of utilities under risk, since we are omitting the classical theory of consumer behavior. When we come to welfare economics, we shall encounter still other interpretations for the class K .

As we have used them so far, the preference-and-indifference relation and the derived preference relation have been nothing but stepping-stones to the definition of the utility function, since the utility function tells us at least as much about the individual's preference field as does the preference relation because of condition (1), (page 9). Nevertheless, we have pointed out that the relation R must satisfy certain conditions in

order that there exist any utility function, u , which satisfies condition (1). One of these conditions on the relation \geq is the requirement of consistency mentioned on page 9. This requirement can be formalized by the condition that \geq must be a weak ordering. In order for R to be a weak ordering, it must satisfy the two conditions (8) and (9).

(8) For all x and y in K , either $x \geq y$ or $y \geq x$;

(9) For all x , y , and z in K , $x \geq y$ and $y \geq z$ imply $x \geq z$.

The condition imposed in (8) states that \geq is a connected relation; in non-formal language, a connected relation is one such that for any two items x and y , either x stands in the given relation to y , or y stands in that relation to x . The requirement of connectedness then simply states that for any two alternatives x and y , either x is preferred or indifferent to y or y is preferred or indifferent to x . Were this condition not to hold, there could be two alternatives such that neither was preferred to the other and which were yet not equally preferable. We would expect that alternatives are comparable, and hence that \geq satisfy the condition of connectedness.

Condition (9) is called the requirement of transitivity, already mentioned above. This condition, too, is one which we would expect to be satisfied by relation \geq .

The definitions (6) and (7) of the indifference and preference relations together with the conditions of connectedness and transitivity logically imply conditions (10) - (14) below. These conditions are listed to show that from the definitions and connectedness and transitivity follow many of the conditions which we would expect $>$, \sim , and \geq to satisfy.

(10) For all x and y in K , exactly one of the following holds:

$$x > y, x \sim y, \text{ or } y > x.$$

(11) For all x and y and z in K ,

$$x > y \text{ and } y > z \text{ imply } x > z.$$

(12) For all x and y in K ,

$$x > y \text{ implies not } y > x.$$

(13) For all x, y , and z in K ,

$$x \sim y \text{ and } y \sim z \text{ imply } x \sim z.$$

(14) For all x and y in K ,

$$x \sim y \text{ implies } y \sim x.$$

The reader can easily verify that these and other conditions which he would expect to be satisfied by \geq , $>$, and \sim do follow from (6), (7), (8), and (9).

Condition (13) is deserving of special attention. This condition implies that for any sequence of alternatives, x_1, x_2, \dots, x_n , such that the relation of indifference holds between any two succeeding pairs, that is,

$$x_i \sim x_{i+1} \quad i = 1, 2, \dots, n-1,$$

it must follow that the first stands in the relation of indifference to the last: $x_1 \sim x_n$. It is easy to imagine a sequence of choices such that we are unable to discriminate between any two succeeding ones, but for which we feel a distinct preference for one of the extremes over the other. This situation is analogous to the case of mass measurements by means of an equal-arm balance. Equality of mass of two bodies is usually operationally defined to mean that the balance remains level when the two bodies are placed in the balance pans. No balance is perfectly sensitive, however, and we may have

a sequence of weights w_1, \dots, w_n each of which differs in weight from the adjacent one by so small an amount as to be undetectable on the balance, but for which the extremes, w_1 and w_n , do differ detectably. We may adopt one of two attitudes towards our theory which will allow us to maintain it in spite of its apparent contradiction of the facts. We may assume that even though the theory does not fit reality exactly, it is a close enough approximation to be useful. Or we may assume that our measuring instruments - the balance and the agent's subjective feelings - are not perfectly accurate.

We have stated above the fundamental condition which the utility function must satisfy (condition (1), page 10); it must reflect the individual's preference-or-indifference relation. This condition is easily restated in terms of the relation \geq :

(1) for all x and y in K , $x \geq y$ if and only if $u(x) \geq u(y)$.

It is a necessary condition for a utility function to exist satisfying (1) that \geq be a weak ordering, as defined by conditions (8) and (9). This is not a sufficient condition, for it can be shown that there are sets K with weak orderings, \geq , for which there exists no function satisfying condition (1). However, we may regard these as pathological cases, and in all the instances we shall be considering, the existence of a utility function is assured if \geq is a weak ordering.

We have noted further that if condition (1) is the only condition placed on u , then u is determined only up to a monotonic transformation. This is a very weak restriction on u , and there would be very little advantage to be derived from working with the utility function rather than the preference-or-indifference relation itself were no more conditions imposed upon u than condition (1). In classical economic theory of consumer behavior, K consists of different commodity bundles which are represented by

vectors $\langle x_1, x_2, \dots, x_n \rangle$ where x_i represents an amount of the i 'th commodity. Then for any given vector, there is a utility, $u(x_1, x_2, \dots, x_n)$. It is usually assumed that u is differentiable with respect to each of its n argument places; that is, that

$$\frac{\partial u}{\partial x_i}$$

exists for $i = 1, 2, \dots, n$. This restriction has some empirical significance since there are relations \succeq for which there are utility functions satisfying condition (1) but none satisfying both condition (1) and the differentiability condition at the same time. Nevertheless, the differentiability restriction serves mainly a conventional purpose in that it helped us to narrow down the class of eligible utility functions, and its empirical significance is generally disregarded. The classical theory of consumer behavior is presented chiefly in the form of differential equations based on a differentiable utility function. We shall not develop this formalism since classical economic theory is not a part of this study.

Another type of restriction on the utility function, which has much more empirical significance than differentiability is one which states some relation between the utility of a combination of alternatives and the utilities of the alternatives of which the combination is composed. We have already encountered one such condition: condition (3) on page 11:


$$(3) \quad u(x_1 * x_2) = u(x_1) + u(x_2).$$

Here x_1 and x_2 are commodity bundles in K , and $x_1 * x_2$ is the bundle which is the sum of the two. This condition has the double function of placing a very strong restriction on the admissible utility function, and at the same time requiring that very strict conditions be satisfied by the relation in order that any function at all exist satisfying (1) and (3). We can illustrate

the empirical significance of this relation in the following way. Let X contain all commodity bundles $\langle x_1, x_2 \rangle$ consisting of just two basic commodities, bread and milk,

and let $u(x_1, x_2)$ be the utility of x_1 loaves of bread and x_2 quarts of milk. Then a combination of two bundles

$\langle x_1, x_2 \rangle$ and $\langle y_1, y_2 \rangle$ is just $\langle x_1 + y_1, x_2 + y_2 \rangle$. By condition (3) then, we must have:



$u(x_1 + y_1, x_2 + y_2) = u(x_1, y_1) + u(x_2, y_2)$. This situation is illustrated in Fig. 1, where the lines of constant utility are shown straight and constant utility differences are represented by lines a constant distance apart. The fact that the lines of constant utility are straight constitutes a restriction on \succeq , since these lines really represent sets of points which are all indifferent to each other. In general, there is no reason to suppose that the set of points representing indifferent commodity bundles should all lie on a straight line, and if they do, it is an empirically significant fact. It can be shown that a necessary condition for a utility function to exist satisfying (1) and (3) is that the sets of indifferent points - called indifference curves - be straight lines. On the other hand, the fact that constant utility differences are represented by lines a constant distance apart does not imply any additional restriction on \succeq because \succeq is completely specified when the indifference curves and the preferences among them are given. This means that it is immaterial which utility is assigned to the

points on one indifference line as long as the utilities assigned for the different lines increase with the distance from the origin. No matter which way utilities are assigned in accordance with the above condition, the corresponding preference-or-indifference relations, as defined by condition (1), will be the same.

For some purposes condition (3) is replaced by (3a) or (3b):

$$(3a) \quad u(x+y) = u(x) + u(y)$$

and

$$(3b) \quad u(x+y) \geq u(x) + u(y).$$

Like condition (3), these also have both an empirical and a conventional significance, i.e., they imply something about \geq , and they serve to restrict the set of utility functions more narrowly than does condition (1).

Condition (3) is called a condition of independence. Thus stated, there is little reason to believe that any ordinary set of alternatives, K , should satisfy it. However, it is frequently of interest to seek to find independent subsets of K , say sets K_1 and K_2 one of which may represent amounts of clothing, which satisfy the condition: for all x_1 in K_1 and all x_2 in K_2 ,

$$u(x_1 + x_2) = f_1(u(x_1)) + f_2(u(x_2))^{(1)}$$

Another way of combining alternatives is in probability distributions. Let x and y be two members of K , and let α be a probability. We define $\langle \alpha x, (1-\alpha)y \rangle$ as a prospect of alternative x with probability α and y with probability $1-\alpha$. For example, suppose x is "go to a movie", y is "study" and $\alpha = \frac{1}{2}$; then $\langle \alpha x, (1-\alpha)y \rangle$ is the prospect of flipping a fair coin to determine whether to go to a movie or to study. If x is "play bridge"

(1) See e.g. Frisch, [11], or Fisher [8]

then

$$x \succ \alpha x, (1-\alpha) y$$

means playing bridge is preferred to taking a 50-50 chance of going to a movie or studying.

In the case of Bernoullian Utilities, with which we are concerned in Part I of this report, the utility of a combination of alternatives according to certain probabilities is simply the expected value of the utilities of the alternatives. In terms of the probability combination operation this condition is formulated: for all x and y in K , and for $0 < \alpha < 1$

$$(15) \quad u(\alpha x, (1-\alpha) y) = \alpha u(x) + (1-\alpha) u(y)$$

Like condition (3) this condition places restrictions on the relation \succ and on the function u .

Finally, we indicate briefly some of the proposed ways of combining utilities or preferences of individuals to obtain a social utility or preference relation. The most obvious method of obtaining a social utility for a group, $S = A_1, \dots, A_n$, of individuals is simply to sum their individual utilities: for all $x \in K$,

$$(16) \quad u(x) = \sum_{i=1}^n u_{A_i}(x).$$

In equation (16), u is the social utility function and u_{A_i} is the utility function of individual A_i . This method of obtaining a social utility, if it is to be meaningful and not an artifact of an arbitrary selection of the individual utility functions from a collection of equally eligible ones, requires that individual utility scales be uniquely determined except for their zero points. If the particular choices of individual utility functions are more arbitrary than simply selecting origins, then the social utility

function obtained by adding individual utility functions selected in one way may differ from that obtained from individual utility functions selected in another way; this difference may be so great that an alternative, x , which is preferred to another alternative, y , according to the first social utility may reverse its relation to y in the second utility function. As none of the conditions so far introduced defines a utility function uniquely except for a choice of zero point, it follows that a social utility function defined according to equation (16) and based on individual utilities defined from these conditions must be arbitrary, and may yield different orderings of the alternatives.

While it is possible to define a social utility from individual utilities in many ways, we would like to require that the social utility functions obtained from the individual utility functions be substantially the same when the individual utilities differ only in arbitrary selection. By 'substantially the same' we mean that the two utility functions should assign the same ordering to the same alternatives.

One way of avoiding the difficulties introduced by the arbitrariness of the utility functions is to define a social preference-or-indifference relation directly in terms of the individual preference-or-indifference relations. If $S = \{A_1, \dots, A_n\}$ is the set of individuals and \geq_i , $i = 1, \dots, n$ are their preference relations, we can represent \geq , the dependence of the social ordering upon them, as follows:

$$(17) \quad \geq = f(\geq_1, \dots, \geq_n).$$

Here f is a function of n argument places whose arguments are relations and whose values are relations. The fact that functional notation is usually associated with functions whose arguments and values are numbers should not confuse the issue, for f is simply a rule assigning to each particular set of individual preference-or-indifference relations a definite social

preference-or-indifference relation (alternatively called a social ordering or social welfare function). The problem of welfare economics as generally posed at present is just what sort of function should be selected. The fact that social preferences are to be based on individual preferences shows that our governing ideals have a utilitarian and democratic basis. But just what this relation should be is still very much in question.

1.3 Questions of Interpretation and Confirmation

There are two fundamentally different ways of interpreting utility theory, each deriving from the use to which the theory is put. The first use is as a descriptive theory about actual individual behavior, one purporting to describe and predict how individuals act in situations of choice. An example of this is the attempt in the theory of consumer behavior to predict the behavior of a large number of individuals and thus the behavior of the market. A second suggested interpretation is that utility theory is a definition of rationality. By this is meant that utility theory does not necessarily describe what an actual person would do in a given situation but states instead that a supremely intelligent person would do in the same situation. The different treatment of the preference ordering of individuals will serve to illustrate the difference between utility as a descriptive theory and utility as a definition of rationality. Under the first interpretation, each person's preference ordering must be transitive to satisfy the axioms. The preference orderings of some individuals, however, might contain circles $x > y$, $y > z$, and $z > x$ which violate the transitivity requirement and seem to be inconsistent sets of preferences. These "inconsistencies" might be explained by the hypothesis that the individual is practically incapable of keeping all his preferences in mind at once and of working out the full implications and logical interrelations among them.

This argument is based on an implicit transformation of utility theory from a descriptive theory to a definition of rational choice making behavior, and explains actual behavior in terms of more or less deviation from the rational norm. These two interpretations of utility might be compared to two possible interpretations of a theory of logic; one as a description of actual thinking processes, the other as a definition of correct thinking processes.

These two interpretations of utility are not unrelated, and it may actually be feasible at times to take the definition of rationality as a good approximation to actuality: It is assumed that each individual tries to be rational, i.e., tries to determine the best means of attaining his desired ends, just as a person tries to think logically, though he may involuntarily fail in both cases. If the choice situation with which the person is confronted is not too complicated, he may be able to think through most of the alternatives and their implications, and arrive at a rational set of preferences, in which case the definition of rationality becomes a descriptive theory.

The third basic interpretation of utility is as a measure of ethical good. The problem here is to determine what, in some sense, is the "best" action for a society or its government to take, given the utility scales of the individuals composing it. Unlike the first two theories, which are concerned primarily with single individuals, the last theory becomes interesting only when the problem is to determine the action of a society of more than one individual. A society composed of a single individual, a "Robinson Crusoe" society, has no ethical problems because it simply acts in accordance with the utility scale of its one member. Under utilitarian ethics, "utility for individual A" and "good for individual A" are identified, and the problem is

somehow to derive a social good, or utility, from the individual utilities, some of which may be in conflict. We shall be concerned only with the first two interpretations throughout the rest of this report.

Before discussing special problems of interpretation and confirmation peculiar to the three different kinds of utility theory, we might point out in a general way how these three types of interpretation affect the problem of confirmation. In utility as a descriptive theory, the problem of confirmation is like that for other scientific theories: if the theory is true, then the statements of the theory must describe actual behavior. Hence it is necessary to compile observations of individual behavior in choice situations and see whether they correspond with what the theory predicts. While this confrontation with experience is, as we shall see, not very straightforward for utility theory, it is at least fairly clear what sorts of tests the theory must meet successfully in order to be acceptable. The ordinary notion of "confirmation" is, however, not applicable to the second kind of utility. These theories are not meant to describe actual behavior, so it is not sensible to test them by confronting them with actual behavior. Interpreted as a definition of rationality, utility theory can only be tested by appealing to a sort of intuitive idea of what rational behavior is like and showing that utility theory does describe this behavior. Many of the arguments used in justification of the axioms of various versions of utility theory make just this kind of appeal. In some cases, the definition of rationality may contain a calculus of utilities by which it is possible to compute utilities for complex alternatives from those of simpler alternatives in a mechanical way. In these cases, utility theory may serve as a mental labor-saving device, much as do the rules of arithmetic which we follow

blindly to save ourselves endless and laborious countings. It is then at least theoretically possible to test the definition of rationality by determining if one's position is actually improved by acting in accordance with it. We shall point out a very simple test of this kind where the definition of rationality includes alternatives with risks.

The same sort of "confirmation" is applicable also to theories of social good based on individual utilities. Obviously it is not sensible to go to actual experience to test a theory of the good, because much of experience is thought to be bad. Hence, once again it is possible to test the theory only by comparing it with our intuitive notions, here, of the "good." By appealing to intuition we are appealing to something vague and possibly contradictory, and it is desirable to try to formulate as precisely as possible just what the "intuitive" conditions are which we expect any theory of social utility to satisfy. It has been one of the great discoveries in welfare economics in recent years that certain of these intuitive conditions are inconsistent. This means that it is impossible to construct a social welfare function out of individual preference fields which satisfies simultaneously all the intuitive conditions for what a good welfare function should be.

We turn now to the problems of interpretation and confirmation peculiar to the different interpretations.

1.3.1 Interpretation and Confirmation in the Descriptive Theory of Utility

We have already indicated some of the alternative interpretations which can be given to the primitive notions, K , \geq , and u of utility theory. We have paid special attention to K , the set of alternatives. Corresponding to each interpretation of K is a variant of each of the three basic types of utility theory.

Only a few of the possible variations in K have so far been indicated. Often a slight change in the interpretation may change the preference pattern radically or even have as a consequence that the axioms are no longer satisfied. We may cite as an example of this a situation which has been described as proof that preference orderings need not be transitive.¹ It has been observed for some animals that they are prone to prefer absence of pain to food and food (absence of hunger) to sex, and sex to absence of pain. If we let "p", "f", and "s" denote pain, food, and sex respectively, the animals' preference relation runs thus: $(\text{not } p) > f$, $f > s$, and $s > (\text{not } p)$. This is not a transitive ordering, and it appears that the behavior of the animals is not described by utility theory. However, we may change the interpretation of K , which had previously included just p , f , and s , to include as well all possible combinations of these three. If absence of pain is preferred to food, this must mean that the combination of no pain, no food, and no sex is preferred to the combination of pain, food, and no sex. In symbols:

$$(\text{not } p, \text{ not } f, \text{ not } s) > (p, f, \text{ not } s).$$

1. Kenneth O. May 36

Similarly, if food is preferred to sex, this could be interpreted to mean the combination of no pain, food, and no sex is preferred to no pain, no food, and sex:

$$(\text{not } p, f, \text{not } s) > (\text{not } p, \text{not } f, s)^{(1)}$$

And finally the third preference may be represented symbolically:

$$(p, \text{not } f, s) > (\text{not } p, \text{not } f, \text{not } s).$$

These three preferences are not circular, and it would in fact be necessary to include all the preferences among all the possible combinations to determine whether or not this relation satisfied the axioms of utility theory. This example should illustrate how critically dependent utility theory is on the interpretation of K .

While it is clear that the sets K and S must be carefully defined, the meaning of "individual" and "alternative" is fairly clear once this has been done. The really formidable problems of interpretation and confirmation arise in connection with the primitive notions \geq and u . We understand \geq intuitively in terms of a subjective feeling of attraction or aversion to the alternatives in K , and imagine that others have similar feelings. However, these subjective feelings are not a good basis for a descriptive theory, since science is in no position to observe them directly, even in cases in which it seems clear that they exist. It is even more problematic to assume that such decision making agents as corporations or governments have feelings, and hence the subjective interpretation must be abandoned entirely in applications of utility theory to this type of "individual." It is, therefore, necessary to look for another interpretation for \geq one which will be scientifically useful.

(1) However, if food is preferred to sex is interpreted to mean:

$$(\text{not } p, f, \text{not } s) > (p, \text{not } f, s),$$

then the preferences are circular. It is not clear from Hay's description which of these interpretations is correct.

Two alternatives are available: (1) to ask the individuals to list the alternatives in order of subjective preference and (2) to deduce the preference pattern from observations of behavior in choice making situations. The first interpretation would seem to be nearer our intuitive idea of \geq , and it would also be more direct. But it must be recalled that utility theory is intended to be a theory of actual decisions, and what people say they would do in a situation is notoriously a very unreliable guide to their behavior when they are actually confronted with it. The second alternative avoids this difficulty, but raises the question of how the preference relation is to be defined from observations of actual decision makings. If x and y are two alternatives, and if x is always chosen over y whenever a choice is presented, then clearly x is preferred to y . The fact that x is always preferred to y demands that the preference-or-indifference relation not change throughout the interval under consideration. Common sense, however, tells us that preference patterns are constantly changing. The fact that \geq may change brings into question the usefulness of utility as a descriptive and predictive theory. Though we may wish to predict the actual choice a person will make when confronted with certain alternatives, utility theory tells us only that he will choose that one with the greatest utility, but not what his preference pattern is, and hence it does not really enable us to predict his behavior. Only under the assumption that \geq does not change does utility become useful predictively. As we know, \geq changes in the ordering it ascribes to certain alternatives: we do not always do the same things under the same circumstances. It is the hope of those who use utility as a predictive theory that the preference-or-indifference relation is relatively stable in the ordering it assigns to the alternatives in which he is interested.

Whether or not \succeq is stable may depend on the definition of K .

For example, suppose K consists of vectors $\langle x_1, x_2 \rangle$ representing x_1 loaves of bread and x_2 quarts of milk. Are x_1 and x_2 to be interpreted as amounts to be acquired in addition to the amount of bread and milk on hand, or as total amounts possessed after the acquisition? Under the first interpretation, $u(\langle x_1, x_2 \rangle)$ is the added utility accruing from the acquisition of x_1, x_2 and it is likely that this will not be stable, but will vary with the amount already on hand. Under the second interpretation, $u(\langle x_1, x_2 \rangle)$ represents the utility of a certain total amount which includes both a new acquisition and what is already on hand, and it may well be that these utilities and the associated preference relation will be fairly stable.

In cases in which \succeq is not stable and K cannot be reinterpreted as indicated above to find a corresponding \succeq which is stable, still another meaning can be assigned to \succeq which does not demand that the same alternatives always be chosen in the same circumstances. In this interpretation, $x \succeq y$ means that the percentage of times, p , that x is chosen when the only alternatives are x and y is greater than or equal to $\frac{1}{2}$. That is $x \succeq y$ means that x is chosen over y , on the average, more than or as often as y is chosen over x . This may be regarded as a generalization of the case of stable preferences, in which it is required that p must be either 1 or 0.

The same remarks as have been made about the old interpretation of \succeq are applicable to the new. It is essential, if utility is to be used predictively, that the relative frequencies of the choices be stable. Here again it is important that K be interpreted appropriately.

I should like to turn now to the problem of confirmation. Even with the revised interpretations of the preference relation, it may be

impossible to test all the statements of utility theory by comparing them with the facts. The relative frequency interpretation of \geq requires that an individual be confronted many times with a choice between alternatives x and y for an estimate of whether $x \geq y$ holds to have a small probability of error. But an individual is very seldom confronted with a choice between just two alternatives. Without the possibility of testing all statements, and particularly, of discovering an individual's preference-or-indifference relation, it becomes a matter of deciding which statements are important for the application under consideration, and trying to test those. In the classical theory of consumer behavior, the aim is to describe the general trends of large masses of buyers. Here it is necessary to assume some uniformity of tastes over large classes in order to generalize from the preferences of the individual. In general, the relation between the axioms of utility theory, or any theory of individual behavior, to the macro-phenomena of social trends is obscure, and it is extremely questionable whether the actual success or failure of the macro-theory is crucially dependent on the details of the individual utility theory on which it appears to be based. The newer theory of Bernoullian utilities, which places more empirical requirements on \geq than the classical theory does, is more susceptible to direct confirmation for this reason. As with the classical theory, though, it is absurd to suppose that its axioms are satisfied exactly, or that preference patterns stay perfectly stable, even when interpreted as relative frequencies. Hence the theory has to be treated as an approximation if it is used predictively at all. The problem of testing whether the theory is a good approximation is difficult, and can only be sensibly attempted relative to certain specified intended applications. As yet, Bernoullian utilities have been exceptionally barren of predictive applications.

1.3.2 The Definition of Rationality

Utility theory as a definition of rationality is concerned mainly with the behavior of single individuals, whether persons or organizations. Here we are no longer concerned with predicting behavior, and so need not require that the preference ordering be stable over time. The object is to describe some of the rules governing the choice to be made among a certain set of alternatives in order for the individual best to achieve his objectives. The theory, then, is to be tested against our idea of what actually constitutes intelligent behavior. We do not expect that the goals of intelligent individuals will always be the same, but only that he should act at any given time to achieve most successfully the goals he has at that time. We do not require either that the things denoted by the primitive terms be objectively observable; it is sufficient for the rational individual to be aware of his own aims at any time, whether or not these are known or knowable to others. The main question to ask of any statement of the theory is always: does it describe rational behavior?

The requirement that \geq be transitive, as part of a definition of rationality, has received some attention recently. The transitivity condition would appear to be an immediate consequence of the transitivity of the ordinary English relation of "better than or as good as." By the rules of English usage, if x is better than or as good as y , and y is better than or as good as z , then x is better than or as good as z . The rational man is supposed to order x and y such that $x \geq y$ if and only if x is better than or as good as y . Therefore, the transitivity of \geq follows. However, this argument really only reflects the original question back to asking why "better than or as good as" should be a transitive relation. In many cases, it is

not at all obvious that rational preference-or-indifference should be transitive, especially where the individual may have several interest which may conflict. May¹ has given us an example in which male students were asked to list preferences among girls as prospective marriage partners, where the girls had various combinations of looks, brains, and no money preferred to plainness, brains, and money, preferred to looks, dullness, and money, preferred to looks, brains, and no money. It can, of course, be claimed that this is an instance of irrationality, but this claim is not easily justified.

Davidson, McKinsey, and Suppel [7] give a more convincing argument for transitivity as follows. Suppose that $x > y$, $y > z$, and $z > x$, and that the individual is presented with a choice among just those alternatives. Then, no matter which one he chooses, there is one he prefers to it, hence he should not have chosen it.

The stipulation that \geq be a weak ordering is also justified to a certain extent by the fact that it simplifies the mathematical problem of treating the preference relations. We have noted that \geq must be a weak ordering in order for a utility function to exist at all, and hence removing this condition removes the possibility of a utility theory.²

1. Kenneth G. May 36

2. An apparent exception to this arises in the theory of consumer behavior. Here the main propositions are expressed in differential equations involving the derivative of the utility function. It may well be that these equations are not integrable, and it has been shown that their non-integrability is equivalent to the intransitivity of the preference relation.

As in the case of the predictive theory, the specification of the class of alternatives is important, and poor choice of K may make it seem that the preference-or-indifference relation is intransitive.

In part I, on Bernoullian utilities, we shall encounter many more assumptions as to what constitutes rationality; these assumptions will be examined there in detail.

1.3.3 The Definition of Good

Most of the remarks made about the problem of confirmation of utility theory as a definition of rationality apply here also. The method of testing can only be a comparison of the statements of the theory with previously held ethical views. As we have pointed out, welfare economics essentially involves interpersonal comparisons of either preferences or utilities. The basic assumption involved lies in the fact that utilities are measures of the good of alternatives to the person involved, and that social good is a function of individual goods.

The question of interpretation may be divided into two categories: (1) which of the set of eligible utility functions actually represents the good? and (2) how shall these utilities be combined to yield a measure of social utility? In connection with question (1), we have pointed out that none of the sets of conditions which the utility function must satisfy defines a utility function uniquely. If the selection of utilities to represent the individuals of the society must be arbitrary, we must require that the social preference ordering defined from these is invariant for arbitrarily differing choices of utilities. If a rule for compounding does not satisfy this requirement, it must be abandoned, or else a new condition must

be sought which restricts the range of admissible individual utilities to such an extent that the social preferences obtained are the same for all equally admissible individual preferences.

The above condition which must be satisfied by any method of compounding individual utilities is a formal restriction which many different methods may satisfy. Presumably the method actually selected will be determined by ethical considerations. We would, for example, probably like to require that the compounding method give equal weight to the utilities of different individuals. We have noted that the problem of defining a social utility invariant under arbitrary changes of individual utility can be bypassed if the social preference scale is defined directly in terms of the individual preference scales. In general, though, the question of what method will be used for this definition is an ethical one, and there is no general agreement on its answer.

In sum we may say that the problems of interpretation and confirmation of this third type of utility theory are similar to those for the theory of rationality, but that the basic ethical principles with which the theory must be compared are much more in doubt than is the intuitive conception of rationality.

2. Bernoullian Utility

2.1. Introduction: the Problem of Rationality

We noted in the Introduction that one of the possible interpretations of utility theory was as a definition of rationality, and noted also the connection which that interpretation has to its interpretation as a descriptive theory. Throughout this part we shall discuss utility theory chiefly as a theory of rationality, because it is in this light that its principles are most easily understood. To appraise a hypothesis introduced as a principle of rational choice it is only necessary for us to review our own intuitive feelings as to whether employment of the principle actually would lead to desirable consequences; whereas to formulate or evaluate the same principle as a description of actual behavior involved us in many complex problems of empirical interpretation and verification. From an empirical point of view, then, we can regard principles of rationality as heuristic guides suggesting empirical hypotheses in fields of behavior in which these are difficult to formulate.

Taken as a theory of rationality, Bernoullian utility attempts to formulate principles of intelligent choice in situations in which the outcome of any choice is subject to chance influences. A simple example serves to illustrate this type of problem. Suppose a man is offered a choice among the following three alternatives: to bet a dollar that an unbiased coin will fall heads, to bet a dollar that the same coin will fall tails (in each case, if he wins, he wins a dollar), or not to bet. We may call these three actions a_1 , a_2 , and a_3 . Besides the set of actions which the man must choose from, there are three possible outcomes: to win a dollar, to break even, or to

lose a dollar.; let us call these x_1 , x_2 , and x_3 , respectively. The man will choose that action which leads to the outcome with greatest utility (assuming he is rational). We observe, however, that not all the actions lead to a certain outcome. Action a_1 is to bet a dollar that the coin falls heads; so under the assumption that the coin is fair, taking a_1 means taking a 50% chance of winning a dollar (if the coin falls heads) and a 50% chance of losing a dollar (if the coin falls tails). Therefore taking a_1 is equivalent to taking a 50% chance of x_1 and a 50% chance of x_3 . In the same way we see that a_2 is equivalent to taking a 50% chance on x_3 and a 50% chance of x_1 , and only a_2 leads to a certain outcome: x_2 (breaking even).

To decide which of the three possible actions to take in the foregoing examples, the man must not only be able to evaluate sure prospects (x_1 , x_2 , and x_3 in this case), but various probabilities of getting there, and Bernoullian utility provides principles of rationality here. It was noted in the general introduction that Bernoullian utility gives a cardinal measure of utility, and it is easy to see from the example why the evaluation of the risk alternatives demands a measure of the relative magnitudes of the values of the outcomes x_1 , x_2 , and x_3 . In trying to decide whether to take a 50-50 chance of winning against losing a dollar, or not to bet, it is not sufficient for the man simply to take into account the fact that he prefers winning a dollar to breaking even, and prefers breaking even to losing a dollar. If the man greatly prefers winning a dollar to breaking even, and only slightly prefers breaking even to losing a dollar, he is likely to risk his money. If, however, he is cautious, and cares less for winning a dollar than for keeping himself from losing a dollar, he will be likely to refuse to bet. In any case, he must take into account the magnitude of his liking for the outcomes, not just the ordinal relationships

(at least, this seems to be true for rational choice).

Before passing to consideration of Bernoullian utility, let us note briefly a distinction which is generally made between theories of decision under risk and decision under uncertainty. The example given above illustrates a problem of decision under risk. In this example the person choosing the action does not know what the actual result of that choice will be: i.e., if he chooses to bet a dollar that the coin will fall heads, the two possible outcomes are winning a dollar and losing a dollar, but at the time of making the decision the man does not know which. He does, however, know the relevant probabilities. In the case of decision making under uncertainty, not only does the man not know what the result of his action will be, but he cannot even assign definite probabilities to the various possible outcomes. We need only change our example of the bet slightly to illustrate the problem of making a decision under uncertainty. Suppose the man is as before required to choose among betting a dollar on heads, a dollar on tails, or not betting. But now, instead of being provided with the information that the coin in question is a fair one, he does not know whether or not the coin is biased, and if it is, what probability it has of falling heads. Still more persuasive examples occur in many familiar situations. Nearly everyone has found himself at some time waiting at a bus stop for a bus about whose schedule he is in almost complete ignorance. It may be late at night, and he does not know whether the last bus has gone yet. He has then to decide whether to wait for the bus or start walking, and if he waits, for how long. This example is as no matter of flipping coins with known probabilities, or even of knowing a definite probability that the bus will come in any interval of time. The information on which the man must base his decision in this case is much less definite.

In the case of decisions under risk, there are intuitively very convincing principles of rational choice (these lead to the construction of the Bernoullian utility function); whereas there are, except in some rather special cases, no such well-founded principles in the case of decision under uncertainty. We shall note some proposals for rationality under uncertainty in the section on applications to theory of games and statistical decisions (section 1.3).

2.2 Bernoullian Utility Functions

2.2.1 The Defining Conditions

If we consider risk combinations of just two outcomes, x and y , and a probability α , then $\langle \alpha x, (1-\alpha)y \rangle$ denotes the uncertain outcome of getting x with probability α and otherwise getting y . Then, if u is a Bernoullian utility function, it must satisfy the conditions:

- (A) $u(x) \geq u(y)$ if and only if $x \geq y$;
- (B) $u(\langle \alpha x, (1-\alpha)y \rangle) = \alpha u(x) + (1-\alpha)u(y)$.

It is easy to see how the fact that u satisfies condition B implies that it must be a cardinal utility. Indeed, if the utilities of any two (not indifferent) alternatives, x_0 and x_1 are chosen, then the utility of any other alternative is determined uniquely by its position in the preference scale. For example, suppose x_0 and x_1 are alternatives of getting nothing and getting \$1.00 respectively and we choose $u(x_0) = 0$ and $u(x_1) = 1$, then the utility of \$2.00, $u(x_2)$, is determined by finding the probability, α , for which the compound alternative of getting \$2.00 with probability α or else getting nothing is held as indifferent to the alternative of getting \$1.00 for sure. If these two alternatives are indifferent, then

$$u(\langle \alpha x_2, (1-\alpha)x_0 \rangle) = u(x_1),$$

and by equation B,

$$\alpha u(x_2) + (1-\alpha) u(x_0) = u(x_1),$$

$$\alpha u(x_2) + (1-\alpha) \cdot 0 = 1$$

$$u(x_2) = \frac{1}{\alpha}.$$

To say, for example, that the utility of any alternative x is twice that of one dollar simply means that a 50-50 chance of getting x or nothing at all is held as indifferent to a certainty of receiving a dollar.

It should be noted that, according to the definition of the Bernoullian utility function, the fact that one alternative may have twice as much utility as another (relative to an arbitrarily chosen zero) says nothing directly about the subjective magnitude of the pleasures due to each. Hence Bernoullian utilities do not necessarily rely on a subjective comparison of magnitudes of pleasure (except as these may enter into the determination of the probabilities at which uncertain alternatives are held as indifferent).

Of course, the mere fact that we have a system of preferences including uncertainties does not guarantee that people do or should hold the utility of an uncertain event to be equal to the expected value of the utilities of the events of which it is composed. More strictly, the existence of a utility function, u , satisfying condition B is not guaranteed for alternatives involving risks. We shall give a set of very plausible axioms from which it is possible to deduce that a Bernoullian utility function exists, but it is worthwhile to note that there is good reason to believe that these axioms are not strictly satisfied, and that other axioms have been proposed for which there exists no Bernoullian utility function. None of the other systems leads to utility functions nearly as simple as the Bernoullian, and

to date none of the alternative systems has been made use of in applications. Abandonment of Bernoullian utilities would have a particularly disastrous effect on the theory of games, in which the equation of condition B is central.

The first explicit use of a Bernoullian utility function was made, as the name suggests, by Daniel Bernoulli, in an attempt to explain the famous St. Petersburg Paradox (which will be described in the section on the utility of money). Bernoulli's hypothesis was that the value of money is not directly proportional to its amount, but rather to its logarithm, and that the value of a risk combination of various amounts is equal to the mathematical expectation of the values of those amounts. Thus, Bernoulli postulated a Bernoullian utility for money, and assumed a special shape for the curve utility vs. money. The fact that Bernoulli suggested that value is not directly proportional to money is what draws attention to his use of utility; however, the combination of utilities according to condition B had been tacitly assumed by theories of gambling before the time of Bernoulli. All these theories were based on the assumption that the gambler's aim should be to follow that course for which the expected value of the money winnings is the most. Even if the value of money is directly proportional to its amount (the rejection of which assumption was Bernoulli's contribution), it is still an additional assumption that these values should combine according to expected values, i.e., according to equation B. The early theories of gambling, therefore, actually assumed a Bernoullian utility also.

The first modern appearance of Bernoullian utilities is in the Theory of games and economic behavior¹ by John von Neumann and Oskar Morgenstern, where it serves as a basis for the theory of games in that the wins and losses are all expressed in utility values. In this book, the assumptions implicit in Bernoullian utility are set forth explicitly for the first time

1. Von Neumann and Morgenstern [19]

in a set of axioms on the preference-or-indifference relation, and it is rigorously demonstrated that a utility function with the required properties, i.e., one which satisfies conditions A and B, exists. Subsequent developments have chiefly taken the form of revisions in the axiom system to get axioms from which the deduction of the existence of the Bernoullian utility function is more transparent, or else suggestions for weakening the system by omitting one or more axioms. The axioms we give first are essentially those of von Neumann and Morgenstern, presented in slightly different notation. Later we shall describe an alternative system, and show how the utility function is constructed from the preference relation.

2.2.2 The Primitive Notions and Their Interpretations

Our axioms (and other related axioms for Bernoullian utilities) are based on three primitive notions: K , \geq , and a "risk operator" which will be explained below. K and \geq have been explained at length in the introduction: K is the class of alternatives, and \geq is the relation of subjective preference or indifference among the alternatives. We have mentioned the risk operator briefly on p. 23 of the Introduction. Let x and y be elements of K , and let α be a probability such that $0 < \alpha < 1$. Then $\langle \alpha x, (1-\alpha)y \rangle$ is a member of K interpreted as the alternative which consists of getting x with probability α , otherwise getting y .

The basic interpreted propositions of the system are just those of the form " $x \geq y$ ", meaning "alternative x is preferred or indifferent to alternative y ." We should note that alternatives are compared only as sure outcomes. Thus to determine whether " $x \geq y$ " holds, we might ask, "If you were given the choice of having x for certain or y for certain,

which would you choose?" It is meaningless to compare a probability α of getting x with a probability β of getting y . Even where a compound alternative, $\langle \alpha x, (1-\alpha)y \rangle$ is compared, the compound is itself regarded as certain, though neither of its components is certain: To say that $\langle \alpha x, (1-\alpha)y \rangle$ is certain is only to say that it is certain that either x or y will occur, but which one is uncertain.

The fact that probabilities appear in our system would appear to involve us in the thorny controversies surrounding the definition of probability. These, however, can be avoided to a certain extent by avoiding interpreting $\langle \alpha x, (1-\alpha)y \rangle$ directly, and seeking to interpret it only in such contexts as $\langle \alpha x, (1-\alpha)y \rangle \geq z$ which are indeed the fundamental ones of utility theory. " $\langle \alpha x, (1-\alpha)y \rangle \geq z$ " means "the compound alternative of x with probability α and otherwise y is preferred or indifferent to z ." We can envisage this compound $\langle \alpha x, (1-\alpha)y \rangle$ as an actual physical alternative, by supposing c is a biased coin with probability α of falling heads. Then $\langle \alpha x, (1-\alpha)y \rangle$ is the alternative which consists of flipping c and taking x if it falls heads and taking y if it falls tails. If utility is interpreted as a descriptive theory, it is not necessary to specify what is meant by " c has probability α of falling heads," since we are interested in predicting the person's behavior which will depend not on actual probabilities but on his subjective estimate of them. Fully expanded, $\langle \alpha x, (1-\alpha)y \rangle \geq z$ means that flipping a coin of subjective probability α of falling heads to determine which of x and y is taken is preferred or indifferent to z . In this case, it seems feasible to omit altogether direct reference to numerical probabilities and express the risk alternatives directly in terms of the component alternatives and the experiment performed to determine which one shall actually take place.

If utility theory is interpreted as a definition of rationality, we may want to interpret the probabilities as objective, or rational, since these presumably represent the subjective probabilities that a rational man would arrive at. However, even here we may still choose to use subjective probabilities (thus by-passing the knotty problem of what objective probabilities are) and arrive at a modified theory of rational behavior, prescribing what a rational man would do, given imperfect estimates of probabilities. Even where probabilities are interpreted as relative frequencies, it is important that the expressions of utility theory are not taken as referring to preferences over a long series of events. The probabilities refer to the particular event (such as flipping the coin) since the choices are defined for particular alternatives. The importance of this will become more apparent in our discussion of the application of utilities to the theory of games.

We do not take as primitive one of the notions discussed in the Introduction, the utility function u . The utility function can be defined uniquely from the preference relation once a zero point and unit of measurement are selected, and hence does not need to be taken as primitive.

2.3 Axioms

The following set of axioms for Bernoullian utilities are in most essentials the same as those of von Neumann and Morgenstern. The principal difference lies in our inclusion of axiom A.3 which states that indifferent alternatives may be substituted for one another to yield indifferent compound risk alternatives. This axiom and axioms A.1 and A.2 insure that as far as the formal statements of the theory are concerned, we may treat the relation of indifference between two alternatives in the same way as logical identity; i.e., we can substitute the name of an indifferent alternative

for that of any alternative in a statement, and the resulting statement is true if the original one is true. The omission of A.3 by von Neumann and Morgenstern indicates that they have taken the interpretation for the class K to be not the set of individual alternatives, but the collection of all sets of indifferent alternatives. We shall not enter into a discussion of this interpretation here, but merely point out that if tacitly assumes some such axiom as A.3, i.e., substitutability of indifferent alternatives.

After stating the axioms, we discuss their significance in terms of the intended interpretation of the primitive notions given above. This discussion will serve both to provide a plausible intuitive justification of the axioms and to indicate some apparent counter-instances of behavior which does not satisfy the axioms. The first pair of axioms - stating that

\succeq is a weak ordering - has been discussed in the Introduction, and this discussion will not be repeated here. After discussing the von Neumann-Morgenstern axioms we shall present some of the alternative axiom sets for Bernoullian utilities.

In the statement of the axioms and in the following discussion, the letters x , y , and z with or without subscripts will denote members of K , and α , β , and γ will denote real numbers in the open interval $(0,1)$; i.e., $0 < \alpha < 1$.

Axioms

A.1 Either $x \geq y$ or $y \geq x$.

A.2 If $x \geq y$ and $y \geq z$ then $x \geq z$.

Definition: $x \sim y$ (x is indifferent to y) for $x \geq y$ and $y \geq x$.

Definition: $x > y$ (x is preferred to y) for $x \geq y$ and not $x \sim y$.

A.3 If $x \sim y$ then $\langle \alpha x, (1-\alpha)z \rangle \sim \langle \alpha y, (1-\alpha)z \rangle$.

A.4 If $y > x$ then $y > \langle \alpha x, (1-\alpha)y \rangle$ and $\langle \alpha x, (1-\alpha)y \rangle > x$

A.5 If $x > y$ and $y > z$, then there exist α and β such that

$$\langle \alpha x, (1-\alpha)z \rangle > y \text{ and } y > \langle \beta x, (1-\beta)z \rangle$$

A.6 $\langle \alpha x, (1-\alpha)y \rangle \sim \langle (1-\alpha)y, \alpha x \rangle$

A.7 $\langle \beta \langle \alpha x, (1-\alpha)y \rangle, (1-\beta)y \rangle \sim \langle \beta \alpha x, (1-\beta \alpha)y \rangle$.

Axiom A.3 says that if x and y are held as indifferent, then the combination of x with probability α and z with probability $1-\alpha$ is indifferent to the same combination with y in place of x . If α is the probability that a certain experiment, c , will succeed, then the outcome of alternatives $\langle \alpha x, (1-\alpha)z \rangle$ and $\langle \alpha y, (1-\alpha)z \rangle$ will be x and y respectively if c succeeds, and z if c fails. In either case, the outcomes are held as indifferent, hence it seems reasonable that the two compounds should be indifferent.

A.3 very clearly rules out an interpretation for the primitive notions under which a preference for a risk alternative, say $\langle \alpha x, (1-\alpha)y \rangle$ is taken to mean that a person prefers to receive commodity x in proportion α and y in proportion $1-\alpha$ in a long series of gambles. As an example, if x is a long playing record, y is 10 ordinary records, we might have x indifferent y , but if z is a long playing record player, probably $\langle .5x, .5z \rangle$ will be preferred to $\langle .5y, .5z \rangle$ if the probabilities are interpreted as relative frequencies, thus violating the axiom. This example is a special

case of complementarity, which we discussed in the Introduction. Our interpretation rules out the possibility that two alternatives x and z , may complement each other by stipulating that at most one can become actual: i.e., not both can actually take place. If we accept the frequency interpretation, we re-introduce the possibility of complementarity by supposing that more than one of a set of alternatives can actually occur, each one happening at some point in a time sequence of events. Perhaps the best way to avoid the possible introduction of complementarity is to adopt the interpretation of K as a set of possible future histories, and then it is clear that at most one can occur.

A₀₄ states that if y is preferred to x then the risk combination of x with probability α and y with probability $1-\alpha$ lies between y and x in the preference scale. The following argument justifies the assumption that $\langle \alpha x, (1-\alpha)y \rangle > x$. The possible outcomes of $\langle \alpha x, (1-\alpha)y \rangle$ are just x and y , and each has a positive probability of occurring, since by hypothesis, y is preferred to x , and of course x is as good as x , hence no matter what happens, the outcome of $\langle \alpha x, (1-\alpha)y \rangle$ is at least as good as x . Moreover, there is a possibility that the outcome will be better than x , since y is preferred to x . Hence we assume that $\langle \alpha x, (1-\alpha)y \rangle$ is preferred to x ; an entirely analogous argument justifies the assumption that y is preferred to $\langle \alpha x, (1-\alpha)y \rangle$.

A₀₅ states that if x is preferred to y and y is preferred to z (i.e., y lies between x and z on the preference scale), then there are probabilities α and β such that $\langle \alpha x, (1-\alpha)z \rangle$ is preferred to y and y is preferred to $\langle \beta x, (1-\beta)z \rangle$. By axiom A₀₄, we know that both $\langle \alpha x, (1-\alpha)z \rangle$ and $\langle \beta x, (1-\beta)z \rangle$ must lie between x and z , so A₀₅ is a kind of continuity axiom which says that whatever y we choose, lying between x and z on the preference scale, there are probability mixtures of x and z

lying on either side of y . Thus, for example, if y lies very close to x , we should expect that the α such that $\langle \alpha x, (1-\alpha)x \rangle$ is preferred to y would have to be close to 1, so that x would be nearly certain to occur.

A.6 asserts that the alternative of getting x with probability α and y with probability $1-\alpha$ is held as indifferent to the prospect of y with probability $1-\alpha$ and x with probability α . The two prospects are actually identical, so the axiom is justified.

A.7 says essentially that the evaluation of a compound alternative $\langle \alpha x, (1-\alpha)y \rangle$ depends solely on the components x and y and the probability of receiving each. The only possible outcomes of the prospect $\langle \beta \langle \alpha x, (1-\alpha)y \rangle, (1-\beta)y \rangle$ are x and y , and the probabilities of their occurring are $\alpha\beta$ and $(1-\alpha\beta)$ respectively. Hence, if the evaluation depends solely on the outcomes and their respective probabilities, then $\langle \beta \langle \alpha x, (1-\alpha)y \rangle, (1-\beta)y \rangle$ must be indifferent to $\langle \alpha\beta x, (1-\alpha\beta)y \rangle$. It is worthwhile to note that the very fact that we have taken $\langle \alpha x, (1-\alpha)y \rangle$ to be a member of the class of alternatives, and thus to have a definite place in the preference scale, represents the tacit assumption that the evaluation of alternatives with risks depends only on the component alternatives, and the probabilities of receiving them. If more than the variables x , y , and α were involved in determining the place of a risk alternative of $\langle \alpha x, (1-\alpha)y \rangle$ in the preference scale, it would be meaningless to talk of 'the alternative' $\langle \alpha x, (1-\alpha)y \rangle$, and its utility.

There are many types of behavior which appear to violate these axioms, some of which we have examined in our discussion of the consistency requirement (which is formalized in axioms A.1 and A.2). The following behavior, which does not seem utterly irrational, contradicts A.3. Let x , y , and z be, respectively, win a dollar, break even, lose a dollar; then a

conservative better might evaluate a .6 probability of winning a dollar against a .4 probability of losing a dollar as indifferent to a certainty of breaking even:

$$\langle .6x, .4z \rangle \sim y.$$

However, he might prefer the prospect $\langle .5y, .5x \rangle$ to $\langle .5\langle .6x, .4z \rangle, .5x \rangle$ (thus violating A.3 which stipulates that these prospects should be indifferent), because the former offers no risk of loss, and a possibility of gain, whereas the latter admits possibilities of both loss and gain. It may be argued that the above described behavior is irrational in that it is difficult to conceive of any particular objective which would be best served by acting in accordance with these preferences. However, this is a negative argument, and in the absence of any clear-cut definition of rationality, it would be impossible to demonstrate conclusively that violating A.3 is irrational.

Marschak¹ has cited the example of mountain climbers for whom it appears that taking a certain small but not zero risk of being killed actually adds to the enjoyment of the climb, and is preferred both to climbs with no risks and to climbs which have a very high risk or certainty of death. The "game" of Russian Roulette affords an even more clear-cut instance of the same kind. In Russian Roulette, the "gambler" is supposed to spin the chamber of a revolver which contains only one cartridge, and without seeing where the chamber stops, to press the muzzle against his hand and pull the trigger. Russian Roulette players evidently prefer to take a chance (1/6 if the revolver is a six-shooter) of being killed to the other available alternatives of not playing at all and having the certainty of not being killed. Presumably they would also prefer to take their chances of Russian Roulette than to accept the certainty of being

killed. If x is the alternative of being killed and y is the alternative of living, and α is the chance that the revolver will fire, then $\langle \alpha x, (1-\alpha)y \rangle$ should represent the alternative of playing Russian Roulette. Then for any player we have $\langle \alpha x, (1-\alpha)y \rangle \succ x$ and $\langle \alpha x, (1-\alpha)y \rangle \succ y$. But this violates axiom A.4, from which it is easy to show that not both of the above two conditions can hold.

Analogous examples, besides the rather bizarre ones given above, of behavior which apparently violates axiom A.4, are easily constructed. As in other instances in which the axioms seem to be violated, it is possible to "explain away" these counter-examples by claiming that we have not chosen the proper interpretation for the set K of alternatives. It seems plausible that the deliberate choice of high risks, as in the case of Russian Roulette, could be explained psychologically on the grounds that the individual desires prestige or attention. If he imagines that attention will be gained by performing the gamble, then he believes that more than just the simple outcomes, x and y , of being killed or staying alive, are involved in the risk combination $\langle \alpha x, (1-\alpha)y \rangle$. Our argument in justification of axiom A.4 relied on the assumption that an outcome, say y , of living, taken as a pure outcome, is the same as the outcome of getting y in the gamble $\langle \alpha x, (1-\alpha)y \rangle$. But our Russian Roulette player evidently believes that the future to be expected if he lives through his gamble is different from the one he would have if he simply lived and avoided the gamble. We may get rid of cases like this by demanding that the mere act of gambling on a combination $\langle \alpha x, (1-\alpha)y \rangle$ cannot alter the outcomes comprising it. However, the price of this "riddance" is very high, since for almost all gambles, the gambling itself affects the thing gambled for; this effect may be indirect, as the fact of having won money

by gambling may influence the attitude of society towards the gambler. It seems best to steer a middle course, not applying the theory to situations in which gambling itself is a prominent element, and hoping that the theory is a close enough approximation to be useful elsewhere. We shall encounter very similar considerations in our discussion of axiom A.7.²

A.5 is a kind of "commensurability" axiom, and we would expect to find contradicting instances in situations in which the alternatives compared are so disparate as to be incommensurable. Let x , y , and z be: respectively, get one penny, get nothing, be executed.² We may suppose that x is preferred to y and y is preferred to z . A.5 then asserts that there is some probability α , such that $\langle \alpha x, (1-\alpha)z \rangle$ is preferred to y , that is, taking a chance α of getting a penny and $(1-\alpha)$ of being executed is preferred to not gambling and getting nothing. But it might very well be that if there were any chance at all of being killed, the person would prefer not to take it just for the possibility of winning a penny. Whether or not such behavior is rational appears to depend on the person's willingness to admit that there exist probabilities for which he would take the risks in question. This is connected with what is perhaps

1. Another type of behavior violating A.4 (and possibly others of these axioms) occurs where the alternative outcomes themselves are action or strategies in a game. For these kinds of alternatives, it is in general impossible to define Bernoullian utilities consistently. This is simply another reminder of the care with which X must be defined. The fact that the set of outcomes, X , cannot contain members which are themselves alternative courses of action in a game makes certain economic applications of the theory of games doubtful, since the objectives in any given economic situation may often be simply to reach economic positions from which the game may be played further to advantage. We shall discuss this in more detail in the section on the decision problem.

2. This example was suggested by analogous examples given in unpublished notes of Raiffa and Thrall.

one of the gravest defects of the system of Bernoullian utilities: namely, that each person is assumed to be able to evaluate risk alternatives for all possible probabilities. Possibly this assumption is justified as a part of a definition of rationality, but it seems absurd as a description of actual behavior. In particular, it is very questionable whether probabilities very close to certainties have much psychological meaning, and these are just the probabilities in question in our example of "incommensurable" alternatives.

Axiom A₀5 is the only one which asserts the existence of an alternative. Therefore, if K is a set of alternatives satisfying all the axioms except A₀5, and K' is a subset of K, a fortiori, K' also satisfies these axioms. In case K does contain incommensurable alternatives, we may choose to isolate a subset, K', of commensurable alternatives; then the subset K' will satisfy all the axioms, and it will be possible to construct a Bernoullian utility function for it. Thus, for some purposes it may be convenient to consider only alternatives roughly comparable to getting a dollar, and leave out such extreme alternatives as being executed. For each commensurable subset of alternatives we obtain a Bernoullian utility function. It has been shown,¹ that the entire set of alternatives, K, can be represented as a set of points in a multi-dimensional vector space in such a way that sets of commensurable alternatives lie on straight lines, and where utility differences for points on the line are proportional to distances between the points.

Finally, there are many instances of behavior violating A₀7. As we would expect, these are instances in which the utility of a risk

1. Hausner and Wendel [35]. What we assert here can easily be deduced from their results, though strictly speaking it is necessary to add the further axiom: A₀3a If $x > y$ then $\langle x, \pi_1(1-\alpha)x \rangle > \langle y, (1-\alpha)x \rangle$ in order to obtain the desired result. A₀3a is deducible from the full set of our axioms, but not from the set from which A₀5 is deleted.

combination depends on more than simply the outcomes involved and the probabilities of receiving each. The typical counter-instance occurs where gambling itself is of value. For example, suppose x and y are winning and losing a dollar, respectively, and α and β are both $\frac{1}{2}$. A.7 asserts in this case that $\langle \frac{1}{2} \langle \frac{1}{2}x, \frac{1}{2}y \rangle, \frac{1}{2}y \rangle$ is indifferent to $\langle \frac{1}{4}x, \frac{3}{4}y \rangle$; that is, that a bet of taking a 50-50 chance of winning or losing a dollar, or else losing a dollar is indifferent to taking a 25-75 chance of winning a dollar against losing a dollar. The actual probabilities of winning a dollar and of losing a dollar are in both cases the same, however, two gambles are involved in the first alternative whereas only one is involved in the second. Then if someone desired the excitement of gambling for its own sake, he might actually prefer the first alternative, and this would contradict the axiom.

The above example is somewhat similar to the case of the mountain climber or the Russian Roulette player, and it might appear that the two axioms A.4 and A.7 stand or fall together because the counter-examples violating them are of the same type. There is, however, a certain difference between the case of the Russian Roulette player and the person who enjoys gambling for its own sake, in that, as we noted, the Russian Roulette player prefers his game because the outcome of living through the gamble is different from and preferable to the alternative of living without gambling, whereas in the case of the man who enjoys gambling, the set of possible outcomes in the gamble $\langle \beta \langle \alpha x, (1-\alpha)y \rangle, (1-\beta)y \rangle$ may be the same as the outcomes in $\langle \alpha \beta x, (1-\alpha \beta)y \rangle$, but the process of obtaining these outcomes is different in the two cases.

2.4 A Different Approach

The following axiom set, similar to one given by Paul A. Samuelson¹ shows a somewhat different formalization of utility theory. The chief point of difference lies in the generalization of the risk operator to include risk combinations of more than two alternatives. This is of course not an essential generalization, since it is always possible to form risk combinations by repeated application of the risk operator on pairs of alternatives. The end result in either case is the same: the derivation of the Bernoullian utility functions. The existence of this function for the two systems shows that they are equivalent, since this implies that both sets of axioms are satisfied. This last assertion is not quite true in the case of the axiom system we are about to propose: we have chosen to impose two arbitrary restrictions which need not in general be satisfied by a system for which a Bernoullian utility exists. These restrictions are not of great conceptual significance, but serve to make the derivation of the utility function particularly easy.

The new system is based on a finite set of "basic" alternatives: A_1, A_2, \dots, A_n . These may be interpreted as a set of mutually exclusive "sure" prospects. From these prospects as a basis, we can construct arbitrary risk combinations as follows: for $i = 1, \dots, n$, $\langle iA_i \rangle$ is a risk prospect interpreted as getting A_i with probability 1. If x_1, \dots, x_k are risk prospects, and $\alpha_1, \dots, \alpha_k$ are probabilities such that $\sum_{i=1}^k \alpha_i = 1$, then

$\langle \alpha_1 x_1, \alpha_2 x_2, \dots, \alpha_k x_k \rangle$ is a risk prospect. We may consider any risk prospect $\langle \alpha_1 x_1, \alpha_2 x_2, \dots, \alpha_k x_k \rangle$ as a ticket in a lottery which awards x_1, \dots, x_k as prizes, with probabilities $\alpha_1, \dots, \alpha_k$ respectively. In

1. Samuelson [27] An important addition to Samuelson's axioms included here was suggested by Professor Howard Raiffa.

general, the lotteries award prizes which are themselves lottery tickets, since for any set of tickets we allow for the existence of a lottery which awards those tickets as prizes.¹ Let K be the total set of lottery tickets of all types built up in this way. If a person chooses any one of the members of K (i.e., any lottery ticket), then he receives whatever prize that lottery yields when run off; if that prize is another lottery ticket, that one too must be run off to determine what the person is to receive. In any case though, what the person finally receives is one of the basic alternatives, since these are the only prizes which are not themselves tickets. It is assumed that from among the entire set of "tickets" the person must choose exactly one, and therefore what he finally receives is just one of the basic alternatives, the one resulting from the particular outcomes of the lotteries involved. Further more, each ticket is associated with a definite probability of receiving the basic alternatives. To each ticket, x , there corresponds a unique associated ticket, x^* , of the following type:

$$\langle .1A_1, \dots, .nA_n \rangle$$

(that is, whose only prizes are basic alternatives), which has the same probability of yielding each of the basic alternatives as does the original ticket.²

Finally, there is a preference-or-indifference relation, defined over the class K . The system thus obtained can be regarded as the

1. We must be careful to prohibit the existence of lotteries which award their own tickets as prizes, or which award as prizes tickets to lotteries which in turn award prizes which are tickets to the first lottery. Such a situation would arise for lottery tickets x, y , and z where $x = \langle .5y, .5z \rangle$ and $y = \langle .5x, .5z \rangle$. If the lotteries are run off in temporal sequence, we must demand that the prizes in a lottery can be either one of the basic alternatives or else tickets to lotteries which are run off later.

2. The "associated" ticket may be defined inductively as follows:

(a) if $x = \langle .1A_1 \rangle$ $i = 1, \dots, n$, then $x^* = \langle .1A_1, \dots, .1A_n \rangle$
(cont'd on p. 59)

generalization of the system involving preferences among risk combinations of just two outcomes. The risk alternative $\alpha x, (1-\alpha)y$ in the von Neumann-Morgenstern system can be interpreted as a lottery ticket to a lottery yielding only prizes x and y . (We assume that the ticket holder gets exactly one of the prizes; in ordinary lotteries it is possible to get no prize, but this can be interpreted as the prize of "getting nothing".) In one respect the present system is less general than the previous one; it starts from a finite set of basic prospects, of which all others are risk combinations, and there is no reason to suppose that this should be the case. The restriction is not fundamental, though, but serves to make the presentation here and the derivation of the utility functions simpler.

The conditions which a Bernoullian utility function must satisfy in our new system are:

- (A) $u(x) \geq u(y)$ if and only if $x \geq y$;
 (B') $u(\langle \alpha_1 x_1, \dots, \alpha_k x_k \rangle) = \alpha_1 u(x_1) + \dots + \alpha_k u(x_k)$.

Clearly, condition B' reduces to condition B (page) in case $k = 2$, since for $\alpha_2 = 1 - \alpha_1$ and $\langle \alpha_1 x_1, \alpha_2 x_2 \rangle = \langle \alpha_1 x_1, (1 - \alpha_1)x_2 \rangle$, hence

$$u(\langle \alpha_1 x_1, (1 - \alpha_1)x_2 \rangle) = \alpha_1 u(x_1) + (1 - \alpha_1) u(x_2)$$

as in condition B.

Using the present set of axioms we shall be able to give a very simple construction of the Bernoullian utility function. In order to make

(cont'd from p. 58) 2. (b) If $x = \langle \beta_1 x_1, \dots, \beta_k x_k \rangle$ and

$$y_1 = \langle \alpha_1 A_1, \dots, \alpha_n A_n \rangle, \quad 1 = 1, \dots, k,$$

$$\text{then } x = \left(\sum_{i=1}^k \beta_i \alpha_{ii} \right) A_1, \dots, \left(\sum_{i=1}^k \beta_i \alpha_{in} \right) A_n.$$

this construction possible we have assumed that there is a finite set of basic alternatives, and that at least two of them are not indifferent (in case all of them are indifferent, then all the risk compounds are indifferent, and the utility function has only one value). There are six axioms:

B.1 \geq establishes a weak ordering over K ; that is

(i) either $x \geq y$ or $y \geq x$

(ii) if $x \geq y$ and $y \geq z$ then $x \geq z$.

As before we define preference ($>$) and indifference (\sim) in terms of \geq .

We further assume that the basic alternatives are arranged in descending order on the scale of preference:

$$A_1 \geq A_2 \geq A_3 \dots A_{n-1} \geq A_n.$$

B.2 $A_1 > A_n$.

B.3 $x \sim x$.

B.4 If $x_i \sim y_i$, $i = 1, \dots, k$, then $\langle x_1 x_1, \dots, x_n x_n \rangle \sim \langle y_1 y_1, \dots, y_n y_n \rangle$

B.5 For $i = 1, \dots, n$, there exists α_i such that $A_i \sim \langle \alpha_i A_1, (1-\alpha_i) A_n \rangle$

B.6 $\langle \alpha A_1, (1-\alpha) A_n \rangle \geq \langle \beta A_1, (1-\beta) A_n \rangle$ if and only if $\alpha \geq \beta$

We shall not attempt to discuss the above set of axioms in the same detail with which we discussed the von Neumann-Morgenstern axioms. Axiom B.1 is, of course, the same as axioms A.1 and A.2 of von Neuman and Morgenstern. B.2 states our assumption that there are at least two alternatives that are not indifferent. Note that A_1 and A_n lie at the extremes of the scale of preferences, with all other alternatives lying between. Axiom B.3 is a sort of generalization of A.7 of the von Neumann-Morgenstern axioms, that is, it asserts that the evaluation of the lottery tickets

depends solely on the basic alternatives which may eventually be received, and the probabilities of receiving each. B.4 is analogous to A.3, a kind of substitutability axiom stating that probability combinations of equivalent alternatives are equivalent. B.5 and B.6 together serve the same function here as was served by A.4 and A.5 in the von Neumann-Morgenstern system, and in fact, A.4 and A.5 can easily be derived from B.5 and B.6. B.5 asserts that for any of the basic alternatives there is some probability combination of the extreme alternatives, A_1 and A_n , which is equivalent to it, and B.6 asserts that in probability combinations of the extreme alternatives, the one which gives the greatest chance to A_1 (the most preferred alternative) is the one that is preferred. Note that if axiom B.2 is not satisfied, i.e., if A_1 and A_n are indifferent, then all probability combinations of A_1 and A_n are indifferent, and B.6 would not be satisfied.

In the following section we shall use this axiom set in constructing the Bernoullian utility function for the class of "lottery tickets".

2.5 The Mapping Theorem; the Construction of the Bernoullian Utility Function; the Uniqueness of the Function

As we have said above, the principal use made of the axioms, as far as applications outside of utility theory itself are concerned, is to derive the existence of the Bernoullian utility function. For example, in the theory of games, utilities are used as the medium in which the payments in the outcomes of the games are expressed, and the preference relation itself does not enter directly at all. It is extremely important then, that the axioms given are sufficient to guarantee the existence of this function, and for this reason, the main derivation from the axioms is the "mapping

theorem¹ which asserts the existence of the Bernoullian utility for preferences satisfying the axioms.

In this part we indicate the construction of the utility function for a preference relation satisfying the axioms of section 2.4.² This construction serves the dual function of showing that a Bernoullian utility function exists, what the empirical significance of the constructed function is. von Neumann has shown³ that a Bernoullian utility function exists for preferences satisfying the von Neumann-Morgenstern axioms. The proof, however, is, quite long, and so we will not even attempt to sketch it here but will confine ourselves merely to stating the result.

Suppose x is any of the compound lottery tickets in the system of section 2.4. According to axiom B.3, x is held as indifferent to its associated ticket \bar{x} , which yields as prizes only the basic alternatives, A_1, \dots, A_n . Suppose $\alpha_1, \dots, \alpha_n$ are, respectively, the probabilities with which A_1, \dots, A_n are received in \bar{x} , i.e.:

$$\bar{x} = \langle \alpha_1 A_1, \alpha_2 A_2, \dots, \alpha_n A_n \rangle$$

By axiom B.5 each of the A_i is indifferent to some combination of the extreme alternatives A_1 and A_n ; therefore there exist probabilities c_i such that:

$$A_i \sim \langle c_i A_1, (1-c_i) A_n \rangle, \quad i = 1, \dots, n.$$

By B.4, the lottery ticket on the right of the above expression can be substituted for A_i in the ticket \bar{x} , and an equivalent ticket results:

$$\bar{x} = \langle \alpha_1 A_1, \alpha_2 A_2, \dots, \alpha_n A_n \rangle \sim \langle \alpha_1 \langle c_1 A_1, (1-c_1) A_n \rangle, \alpha_2 \langle c_2 A_1, (1-c_2) A_n \rangle, \dots, \alpha_n \langle c_n A_1, (1-c_n) A_n \rangle \rangle$$

1. The mapping theorem is so called because it asserts that the alternatives can be mapped onto the real numbers (via the utility function) in such a way that certain important relationships among alternatives are reflected in parallel relationships among the corresponding numbers.

2. This construction was suggested by Professor Howard Raiffa.

3. von Neumann and Morgenstern, [9] Appendix.

Let y be defined:

$$y = \langle \alpha_1 \langle \epsilon_1 A_1, (1-\epsilon_1) A_n \rangle, \alpha_2 \langle \epsilon_2 A_1, (1-\epsilon_2) A_n \rangle, \dots, \alpha_n \langle \epsilon_n A_1, (1-\epsilon_n) A_n \rangle \rangle$$

' y ' is a compound ticket (a ticket whose prizes are the tickets $\langle \epsilon_i A_1, (1-\epsilon_i) A_n \rangle$), and hence has an equivalent associated ticket, \bar{y} . y is the ticket which yields as prizes only the basic outcomes which can be received by playing out y , with the probabilities for these alternatives as given by y itself. The only possible outcomes of y are simply A_1 , and A_n , and the probabilities of receiving these are, respectively,

$$\sum_{i=1}^n \alpha_i \epsilon_i,$$

and

$$\sum_{i=1}^n \alpha_i (1-\epsilon_i) = 1 - \sum_{i=1}^n \alpha_i \epsilon_i$$

Therefore

$$\bar{y} = \langle (\sum_{i=1}^n \alpha_i \epsilon_i) A_1, (1 - \sum_{i=1}^n \alpha_i \epsilon_i) A_n \rangle$$

Since $x \sim \bar{x}$, $\bar{x} \sim y$, and $y \sim \bar{y}$, then $x \sim \bar{y}$. Hence for all x there is at least one lottery ticket of form y (i.e., an associated ticket which yields as prizes only the basic alternatives A_1 and A_n) to which x is indifferent. It is an easy consequence of axiom B.6 that there is at most one ticket of form y equivalent to x , so x is equivalent to a unique ticket of the form

$$\langle \alpha A_1, (1-\alpha) A_n \rangle$$

Let x^* be this ticket. We can now construct the utility function directly in terms of the ticket x^* . That is, if x is any ticket, and

$$x^* = \langle \alpha A_1, (1-\alpha) A_n \rangle$$

we define

$$u(x) = \alpha.$$

The above two equations imply:

$$x \sim x^* = \langle u(x) A_1, (1-u(x)) A_n \rangle$$

Similarly, if y is any other lottery ticket,

$$y \sim \langle u(y) A_1, (1-u(y)) A_n \rangle$$

and by axiom B.6

$$\langle u(x) A_1, (1-u(x)) A_n \rangle \geq \langle u(y) A_1, (1-u(y)) A_n \rangle$$

if and only if $u(x) \geq u(y)$. Therefore $x \geq y$ if and only if $u(x) \geq u(y)$,

and the function, u , satisfies condition A. To show that u is a Bernoullian utility it is only necessary to show that it satisfies condition B'. Con-

sider the ticket $x = \langle \alpha_1 x_1, \dots, \alpha_k x_k \rangle$; by the definition of u ,

$$x_i \sim \langle u(x_i) A_1, (1-u(x_i)) A_n \rangle$$

Therefore, according to axiom B.4,

$$\langle \alpha_1 x_1, \dots, \alpha_k x_k \rangle \sim \langle \alpha_1 \langle u(x_1) A_1, (1-u(x_1)) A_n \rangle, \dots, \alpha_k \langle u(x_k) A_1, (1-u(x_k)) A_n \rangle \rangle$$

But the associated ticket of the compound ticket on the right above is:

$$\langle \left(\sum_{i=1}^k \alpha_i u(x_i) \right) A_1, \left(1 - \sum_{i=1}^k \alpha_i u(x_i) \right) A_n \rangle$$

By definition of $u(\langle \alpha_1 x_1, \dots, \alpha_k x_k \rangle)$,

$$\langle \alpha_1 x_1, \dots, \alpha_k x_k \rangle \sim \langle u(\langle \alpha_1 x_1, \dots, \alpha_k x_k \rangle) A_1, (1-u(\langle \alpha_1 x_1, \dots, \alpha_k x_k \rangle)) A_n \rangle$$

hence

$$u(\langle \alpha_1 x_1, \dots, \alpha_k x_k \rangle) = \sum_{i=1}^k \alpha_i u(x_i),$$

Hence condition B' is satisfied, and u is a Bernoullian utility function.

We are now in a position to understand the significance of the utility function thus constructed. We have seen that for any compound alternative, x , x is indifferent to the lottery ticket, $\langle u(x)A_1, (1-u(x))A_n \rangle$. Hence, $u(x)$ is simply the probability such that the alternative A_1 with probability $u(x)$ and A_n with probability $(1-u(x))$ is indifferent to x . Since $u(x)$ is a probability, it must lie in the interval $[0,1]$; i.e. $0 \leq u(x) \leq 1$. This may seem surprising, but is easily explained when we consider that we have taken the worst alternative, A_n , as the zero point of the scale, and the best, A_1 , as the unit point, and that all other alternatives lie somewhere between A_1 and A_n on the scale of preferences and therefore between them in utility value. To take a concrete example, suppose that A_1, \dots, A_n are alternatives of receiving incomes of different amounts ranging from a million dollars a year down to minus a million (if that is possible), on the average, for life. These can be ranged in dollar increments, and A_1 is the prospect of receiving a million a year and A_n ($n = 2,000,001$) is the prospect of losing a million. Then we arbitrarily select A_n to have utility 0 and A_1 to have utility 1, and we expect all other alternatives to have utilities somewhere between 0 and 1. To determine the exact utility of a prospect x , we simply determine the probability $u(x)$, of which we would believe that a chance $u(x)$ of getting a million a year for life and $1-u(x)$ of losing a million a year for life is just an even trade with x itself.

The arbitrary selection of a zero point and a unit point in the utility scale is entirely consistent with the conditions A and B (or B'), defining the utility function. It can be shown that if u satisfies these conditions, then any other function, u' , related to u by the equation

$$u'(x) = au(x) + b,$$

(where a and b are any real numbers, such that $a > 0$), is also a Bernoullian utility function. This implies that (at least as far as conditions A and B are concerned), the choice of the particular utility function is arbitrary to the extent of the selection of zero and unit points. Once these points are determined, however, the utility function is fixed uniquely, as can be seen from our construction of the utility function for the system of section 2.4, in which choosing A_n and A_1 as zero and unit points is accompanied by a unique determination of the utilities of the other alternatives.

3. The Decision Problem

3.1 Description of the Decision Problem and its Relationship to Utility Theory

The decision problem to which we refer in this section is simply the problem of deciding, in what is in some sense the "best" way, among a number of alternative courses of action. A theory of decision procedures, like utility theory, may be interpreted as either a definition of rationality, or as an empirical theory purporting to describe actual human behavior. Utility, in particular Bernoullian utility, is said to be a special case of a decision theory, because it relates to choices (i.e., decisions), in situations in which the outcome of the selected action involves risk. As a definition of rationality, Bernoullian utility provides certain principles or precepts which it can be plausibly argued a rational person should follow in making decisions among alternatives involving risk. If the alternatives under consideration involve no more than simple calculable probabilities (such as are exemplified by alternatives whose outcomes are dependent on events like the fall of a coin or die, or a lottery), it is hard to imagine any more rules of rational choice beyond those implied by the axioms of Bernoullian utility. That is, given a set of basic preferences and intensities of preferences, which we assume are entirely arbitrary and hence not prescribed by rational rules,¹ the axioms of Bernoullian utility fully prescribe the preferences among the risk combinations of these alternatives, hence completely solve the decision problem for choices among alternatives which involve risk² only.

1. Except for the requirement that the preference ordering be "consistent".

2. See 1.1 for discussion of the meaning of "risk".

Once we get beyond simple risk combinations of alternatives, we are once more confronted with the need for principles (either prescriptive or descriptive) of choice among these enlarged sets of alternatives. Let us examine some of these new kinds of decision situations. We noted in the introduction that some choices involve uncertainty, as opposed to risk. As a typical example of a decision under uncertainty there is the case of the man trying to decide between walking home or waiting for a bus late at night when he does not know what the bus's schedule is and whether the last one has gone. Within the realm of practical problems of this type which it is important to try to handle systematically, are statistical decision problems. Such a problem arises when a firm attempts to decide on the basis of some sample data whether to accept or reject a certain lot of goods. In this case, the actual composition of the lot (the percentage of defective items within it) is unknown, and the company aims to accept only lots which meet its standard and reject only those that do not. The company then has to decide using the limited information provided by the sample. The uncertainty that is involved here lies in the fact that the company does not know at the time it makes its decision whether it made the right one, or even, in most cases, what the probability of an error is. Rather than attempt to give any clearer exposition of statistical decision problems here, we shall defer extended discussion of them to Part 3.3.

Still another kind of decision is involved when making moves in a game. The central feature of these types of decisions is that the final outcome of any course of action depends not only on the action, or on either known or unknown random factors, but also on the actions of an opponent who may be rational and try to anticipate the other's action in order to turn it to his own advantage. Hence we arrive at the theory of games, the master

for whom modern Bernoullian utility was created as servant. As we have presented the problem of decision under uncertainty, decisions in game situations are simply special cases, in that the result of an action in a game is uncertain to the extent that it depends on the opponent's act as well. Strictly speaking, the term "uncertain" is usually reserved for alternatives whose outcome is completely determined by the act chosen (i.e., the "decision"), plus certain random events, of which the relevant probabilities are not known, and are not dependent on the actions of a calculating opposition. The difference between decision in game situations and decisions under uncertainty lies in the fact that the outcome of an action in a game depends on acts of players whose own actions depend in turn on the first player's actions; whereas the extrinsic factors affecting the outcome of a decision under uncertainty are independent of the decision made.¹ For example, in the question of whether or not to wait for the bus, and if so, for how long, the outcome depends on the man's decision and on the actual schedule followed by the bus. However, the bus's schedule is independent of the man's decision to wait, so what we have is a case of uncertainty. This could be easily changed into a game situation, however, by making up a malevolent schedule which directs that the bus avoid, as much as possible, picking this man up.

1. Some of the distinctions made here and elsewhere stand in considerable philosophic doubt. We refer particularly to the notion of "independence" used here, and also to the distinction between risk and uncertainty, depending, as it seems, on the distinction between "known" and "unknown" probabilities, and on the notion of "randomness." While we do not agree with those who would relegate the activity of clarification of these concepts to the limbo of useless philosophizing (indeed, it has been just such queries which have been at the root of the many reformulations of statistics and game theory of recent years), we believe this is not the place to attempt such investigations.

Having thus emphasized the difference between game decisions and decisions under uncertainty, it remains to point out that the two problems are sufficiently similar that it is possible to go some distance in treating them before it is necessary to make the distinction. Even then the distinction may mislead by inducing one to believe that decision problems are definitely of one kind or the other, whereas the fact is that many problems involve elements both of uncertainty and of a game.

We have said that Bernoullian utility theory "solves" the decision problem in the special case in which the final outcomes depend only on the action chosen, and on random factors for which the probabilities are known. The solution consists in the fact that each of the alternative actions is equivalent to selecting some probability combination of the final outcomes (analogous to choosing a lottery ticket for which the final outcomes are the prizes), and hence it is possible to assign a utility to the actions themselves which is completely determined by the utilities of the consequences of the actions and the probabilities with which they occur. In deciding on a course of action, then, one simply selects that with the highest utility. In passing to the general decision problem, we still have a set of possible actions and a set of final outcomes, and it is useful to represent the final outcomes in terms of their associated Bernoullian utilities. Thus in the theory of games, it is useful to represent the possible rewards from the game in terms of their associated utilities. In general, the object, as in the case of alternatives under risk, is to determine which of the available actions yields the best outcome as measured in utility. In the case of alternatives under risk, the best action was that one which itself had the highest Bernoullian utility (i.e., yielded greatest expected utility of outcome). As we shall see, it is not possible to assign

Bernoullian utilities directly to the alternative actions in the general decision problem, hence it is not possible to solve it by selecting the action with greatest utility. Nevertheless, the representation of the outcomes as utility payments has many advantages. For one thing, all the final payments are reduced to comparable quantities, whereas the actual physical situation may involve payments of very disparate kinds, such as social prestige, money, foodstuffs, etc. Secondly, chance events may be included among the possible outcomes, since Bernoullian utilities are defined for these as well as for certainties. The simple relation between the utility of a risk alternative and the utilities of the sure alternatives of which it is compounded has profound consequences for all of the theories of decision so far developed.

3.2 Formalization of the Decision Problem¹

In accordance with the terminology of Gershick and Blackwell, we shall speak of all decision situations as games, though in fact any given situation may involve no competitive factors at all. Each "game" will involve a number of "players" whom we call $1, 2, \dots, n$. It is to be understood that the first player is the one in whose decision we are interested, and the other "players", $2, \dots, n$ are his "opponents". The terminology of games is used only for convenience and uniformity. In general, the players, especially the opponents, need not be human beings,

1. Much of this material is drawn from Gershick, H.A., and D. Blackwell, Theory of Games and Statistical Decisions, to which we are deeply indebted. The notation employed here differs considerably from theirs.

and may be any factors which influence the outcome in a decision situation. Associated with each player, $i = 1, 2, \dots, n$, is a set of possible "actions" S_i , which we will call the strategy space of player i . Each player is assumed to make his decision in complete ignorance of what choices the opposing players have made. When each player has moved, the outcome is determined. We denote the outcome resulting from moves s_1, \dots, s_n by $E(s_1, \dots, s_n)$. The set of possible outcomes, then, is the set of events, $E(s_1, \dots, s_n)$ for every possible combination of actions by the players. In conformity with our former notation, we denote the set of outcomes by K .

For each outcome, k in K , there is a set of "payoffs" to the players. Instead of working with the payoffs directly, we shall work with Bernoullian utilities associated with them. The utility to player i from outcome k will be denoted $u_i(k)$. Then for each set of actions, s_1, \dots, s_n by the players, there is a corresponding payoff in utilities $u_i(E(s_1, \dots, s_n))$. To avoid the cumbersome notation $u_i(E(s_1, \dots, s_n))$, we let the function $M_i(s_1, \dots, s_n)$ represent the utility payoff to player i , i.e.,

$$M_i(s_1, \dots, s_n) = u_i(E(s_1, \dots, s_n)).$$

M_i will be called the "payoff function" for player i .

In some cases it is inappropriate to speak of utility payments to some of the players, since in fact they may represent inanimate factors in the decision situation, and hence have no preference. In these cases we will assume that u_i and M_i are undefined for the "player" i in question. In the case of statistical decisions, the only animate player is player 1, and his opponent is nature. In this case we designate player 1's utility and payoff functions u and M , omitting the subscripts.

Whether or not the opposing players are animate, it is player 1's problem to choose from the set S_1 of actions available to him that action,

s, which will yield him the best outcome. This is the decision problem.

Let us apply the abstract scheme just set up to some concrete situations. A special, trivial case is the "game" involving only one player with possible action x and payoff function M . For each act x he receives utility $M(x)$, and hence his decision problem is solved by selecting that action for which the utility payment is greatest.

The most obvious application is to games. The simplest game, aside from the 1-person game, involves but one move by each of the players, after which the payoff is made. In this case, our scheme can be applied directly. The moves available to any player $i = 1, \dots, n$ are just those of the set S_i and $M_i(s_1, \dots, s_n)$ is player i 's payoff for moves s_1, \dots, s_n . Games of a more complicated nature, involving several moves, can be reduced to the scheme through the use of what are called "strategies." Suppose, for the sake of simplicity, we are considering the 2-person game tic-tac-toe, in which player 1 has the first move. A strategy for player 1 in this case is some rule or set of rules which tells him unequivocally what move to make in every situation in which he may find himself. A strategy for player 2 is defined in the same way. Then clearly if the two players each pick a strategy in advance, there is no need for them to play the game out, because they could both tell their strategies to a referee who could carry out the indicated moves and determine the winner. The game is now reduced to a single move by each player, namely, picking a strategy from among the strategies which are permissible according to the rules of the game.

If we call the choosing of the strategy by a player his "move", we have reduced the game to a one-move form. In the general game situation, the sets S_i (which we have already designated "strategy spaces"), are just the available strategies to player i , and the game consists of just one move by each player, the act of choosing a strategy.

As yet we have not mentioned games which involve chance moves, though these are indeed the great majority of games we know. Any game of cards, for example, involves the chance factors which determine the order of cards in the shuffled deck and therefore which cards are dealt to whom. A strategy in poker tells the player just what to bet (or whether to withdraw) and if there is any leeway in handling his cards, how to do that. But even if all the players have determined on certain strategies, the outcome of an actual play of a game is not determined completely, because this depends also on what cards are actually dealt. One way to treat this situation is to introduce another player, player $n + 1$, with available strategies Ω , whom we might call "nature" and who determines the outcome of all chance occurrences. In the case of poker, the chance occurrences are confined to the shufflings of the deck, so the space of nature's strategies, Ω , is just the set of possible shufflings. Then for each choice of strategy s_i by the players, and ordering of the cards by nature, ω , there is a uniquely determined outcome, $E(s_1, \dots, s_n, \omega)$ and payoff to the players $H_i(s_1, \dots, s_n, \omega)$.

For many reasons it is not desirable to include another player, "nature", whose moves are, so to speak, blind, among the set of actively competing, self-interested players. The actual solution to a particular decision problem depends not only on the strategies available to the players, but also on their motivations for playing, represented by the payoff functions. The inclusion of a player who acts randomly, without motivation, means that a situation involving this type of player cannot be analyzed in the same way as one involving only intelligently competing players. Our second method of representing games with random moves avoids the introduction of this additional player. Let us for a moment imagine the $n + 1$ st

player still included in our n -person poker game. Then for strategies s_1, \dots, s_n for the regular players, and shuffle W , there is a unique outcome $E(s_1, \dots, s_n, W)$. Let us suppose that there are only a finite number of strategies available to nature (as is the case for shuffles of a deck of cards); we label these W_1, W_2, \dots, W_m . Each of these has a definite probability¹ $\alpha_1, \dots, \alpha_m$, since nature's moves are random. For each choice of strategy by the players s_1, \dots, s_n we can define the "risk outcome" $E'(s_1, \dots, s_n)$ where for all $j = 1, \dots, m$

$$E'(s_1, \dots, s_n) = E(s_1, \dots, s_n, W_j)$$

if W_j occurs. Thus $E'(s_1, \dots, s_n)$ is the risk alternative

$$\langle \alpha_1 E(s_1, \dots, s_n, W_1), \alpha_2 E(s_1, \dots, s_n, W_2), \dots, \alpha_m E(s_1, \dots, s_n, W_m) \rangle$$

in the notation of part 2.4. In this way, an extended set of outcomes, K' , is defined, including risk outcomes from the former set of outcomes K .

Bernoullian utilities are defined for risk outcomes, hence we can define new payoff functions M'_i for the players, and the new functions are related to the original functions M_i by the equation²

$$M'_i(s_1, \dots, s_n) = \sum_{j=1}^m \alpha_j M_i(s_1, \dots, s_n, W_j).$$

1. In general, these probabilities $\alpha_1, \dots, \alpha_m$ need not be independent of the choice of strategies by the other players. The α 's themselves may be functions of s_1, \dots, s_n . What is crucial is that the players know the probabilities, and how they depend on their own choices of strategy. In case the players do not know the probabilities, it is not possible to carry out the reduction here outlined.

2. This follows immediately from condition B', p.[59].

We have applied the formalization so far only to decisions in games; the same scheme can, however, be used to represent decisions in other than game situations. We have mentioned the decision problem under uncertainty. This can be formulated as a "game" with two players: player one, who attempts to make the decision, and player two, nature, who may be in one of a number of "states" about which player one is ignorant. The actions or "strategies" available to player one and nature are, respectively, S and Ω , and for each choice of strategies s and ω by the players, there is a payoff $H(s, \omega)$ in utilities to player one. It makes no sense to talk of a payoff to nature, since it is assumed that nature is indifferent to the outcome. To represent the problem of the man trying to decide how long to wait for the bus before he gives up and walks home, the man's strategies are just the set of time intervals he could wait before starting to walk, and nature's strategies are just the different possible times at which a bus might leave the stop where the man is waiting. Once man and nature have chosen their strategies, the outcome is certain: if the waiting interval chosen by the man is such that a bus arrives in it, the man rides; otherwise he walks. Therefore, there is a definite utility for the man, $H(s, \omega)$ associated with each pair of strategies s and ω .

Unlike the case in which nature's choice of strategy is made according to known probabilities, as in games of chance, it is not possible to suppress nature's role in the decision "game" by introducing risk outcomes in the case of uncertainty. In the bus example, the man will in general not be able realistically to assign definite probabilities for the arrival of a bus in any interval of waiting. If the man did know, for any possible interval of waiting, s , a probability p_s that the bus would arrive in that interval, then the problem could be simplified as follows. Let ω_s

be the strategy of nature of having a bus arrive in the interval s , and \bar{U} be the strategy of not arriving in that interval. Then for each interval, s , there is a definite risk outcome

$$\langle p_s E(s, U_s), (1-p_s)E(s, \bar{U}_s) \rangle$$

That is, a chance p_s that the bus arrives during s , and $(1-p_s)$ that it does not arrive. If $H(s, U)$ is the utility of strategies s and U , we can define a payoff function M^s for the strategy s alone by the equation

$$M^s(s) = p_s H(s, U_s) + (1-p_s)H(s, \bar{U}_s).$$

$M^s(z)$ is simply the Bernoullian utility of the risk alternative given above. From what we have shown, if it is possible to define probabilities for the times of arrival of the bus, then we can reduce the problem to a 1-person game for which the decision problem has a trivial solution.

3.3 Mixed Strategies

We have said that the sets of actions or strategies available to players $1, \dots, n$ are s_1, \dots, s_n . From any given set of actions, s_1 , available to player 1, we can generate a larger set of actions in the following way. Suppose s_1 is finite, and consists of the strategies s_1, \dots, s_m . Then instead of choosing one of the strategies outright, player 1 can let some chance device decide which strategy he will use. The chance device must then give a certain definite probability σ_i to each strategy s_i , where

$$\sum_{i=1}^m \sigma_i = 1$$

since exactly one strategy must be chosen. The act of allowing the strategy to be chosen by a random device is called a mixed strategy. Each mixed strategy is represented by a distribution function, σ , where for $i = 1, \dots, m$, (s_i) is the probability with which strategy s_i will be chosen. To be able

to distinguish the original set of strategies from the mixed strategies, we will refer to the original strategies as pure strategies. The set of possible mixed strategies for player 1 is just the set of distribution functions for the strategies s_1, \dots, s_n . We denote the set of mixed strategies for player 1, Σ_1 . We now envision an enlarged game of n players in which each person's move consists of picking a mixed strategy from among those available to him.

Once these strategies are picked, the game's outcome is determined completely except for random factors. The random factors may enter at two points: first, in the operation of the chance devices which determine which of the pure strategies the players are to follow, and second, at random moves within the play of the game. Since the probabilities involved in the chosen mixed strategies are known, once the strategies are picked a definite probability can be assigned to each possible outcome of the play. Hence a choice of mixed strategies is associated with a risk outcome. To illustrate in the case of a 2-person game, suppose $\{s_1, \dots, s_n\}$ and $\{t_1, \dots, t_n\}$ are the strategy spaces of players 1 and 2 respectively. For any choice of pure strategies s_i and t_j by the players, there is a unique outcome $E(s_i, t_j)$. Let player 1 now choose to play according to mixed strategy σ (i.e., to choose s_i with probability $\sigma(s_i)$) and player 2 to play according to the mixed strategy τ . This choice itself determines the risk outcome $E^i(\sigma, \tau)$ which is for $i = 1, \dots, n$, $j = 1, \dots, n$ to get outcome $E(s_i, t_j)$ with probability $\sigma(s_i) (t_j)$. The risk outcome is associated with a Bernoullian utility, and a payoff function M^i for the mixed strategies satisfies the condition:

$$M(\sigma, \tau) = u_1(E^1(\sigma, \tau)) = \sum_{i=1}^n \sum_{j=1}^n \sigma(s_i) \tau(t_j) M(s_i, t_j).$$

The reader can easily generalize this to the case of n players. Henceforth, we shall use M_i to denote the payoff function to player i both for pure strategies and mixed strategies.

It may be asked why mixed strategies should be considered among the possible strategies or actions which a person might take in a decision situation. Since a mixed strategy is only taking a chance among a number of pure strategies, and a pure strategy must eventually be followed anyway, it might be thought that a mixed strategy can be no better than the best pure strategy of those among which the mixed strategy selects. If, for example, the mixed strategy picks one of two pure strategies with equal probabilities, then it must pick one of them, and it might be thought that the best it can do is pick the best pure strategy.¹ Then why not pick this in the first place? To answer this satisfactorily we shall have to go more deeply into the theory of games and the concept of a solution to a game. We can only hint here that the concept of a "best" strategy is not clearly defined, and that in games against an intelligent opponent, choosing a mixed strategy has the effect of making it impossible for him to predict what action will be taken, and hence what effective counter-measures are required.

3.4 Strategies in Statistical Games

Statistical decision problems can be represented as games between two players: player 1, the statistician, and player 2, Nature. The statistician is assumed to have definite preferences as to the outcome, and Nature

1. This is precisely the heuristic argument given for axiom A.4, p. 49. The fact that this argument breaks down here, and the reasons why, suggest important restrictions on the range of application of Bernoullian utilities.

is assumed to be indifferent. Statistical decision problems are distinguished from general problems of decision under uncertainty by the fact that the statistician can obtain some information about nature by experimenting or sampling before he makes his decision. Let us assume that the statistician has a definite set of final actions, A , among which he must eventually choose, and that the set of possible states which Nature can be in is Ψ . Given a choice of action, a in A , by the statistician, and a state φ in Ψ of Nature, there is a definite outcome, $E(a, \varphi)$ which is associated with a payoff in utility $U(a, \varphi)$ to the statistician. We must be careful in the case of statistical games to distinguish between the sets of actions and states, and the strategy spaces for the statistician and Nature respectively; the strategy spaces for the two players are defined in terms of the basic sets of actions, but also involve the possible experiments which the statistician can perform.

Besides the final actions allowed the statistician, he is also allowed to experiment to obtain information about Nature. We may divide decisions based on experiment into two categories: if the nature of the experiment to be performed is completely determined in advance, as far as the physical operations performed and observations made are concerned, the corresponding decision problem is called a single experiment game; if the experimental operations are not determined in advance, but may vary according to what the previous observations in the experiment have been, the decision situation is called a sequential game.

We may illustrate the two types of experiment by a single statistical decision problem. Let us suppose that a manufacturer of electrical equipment has received a shipment of 1000 fuses, which he suspects may contain so high a proportion of defectives that it would be more profitable

to return the lot to the shipper than to use them in his equipment. He may decide to sample the lot by selecting 20 fuses at random and testing them to find out how many defectives there are in the sample, and then base his decision on the result. This type of sampling procedure is an example of a single experiment decision.¹ Here the operations to be performed are all specified in advance: they are simply to select and test 20 fuses, recording the number of defectives. The manager might have performed the following type of experiment instead: to select fuses one at a time at random from the lot, testing each one as it is selected, and noting whether or not it is defective, and stopping after either (1) twenty fuses have been tested or (2) a total of eleven defective fuses have been found. He might, for example, decide to accept the lot if fewer than eleven defectives are found in the first twenty, and reject it otherwise. In any event, this test represents a sequential experiment, since the actual operations performed in carrying out the test are not completely specified in advance, and may, in fact, vary all the way from testing a minimum of eleven fuses to testing a maximum of twenty.

The set of outcomes of the experiment is called the sample space, and it is in terms of the sample space that the strategy space for the statistician is defined. The statistician bases his decision on the outcome of the experiment, and a complete solution to the decision problem requires him to decide in advance what action he will take in case any one of the possible outcomes of the experiment is observed. A strategy for the statistician is a "complete" decision of this type: formally, a strategy is a

1. No confusion should result from the fact that this "single" experiment can be regarded as a succession of 20 smaller experiments; the important point to be kept in mind is that the sequence of operations is predetermined.

function, d , defined over the sample space (which will be denoted Z), and taking values in the set of actions, A . For all z in Z , $d(z)$ is a member of A , and if d is the strategy decided on by the statistician, and if z is the actual outcome of the experiment, then $d(z)$ is the action taken.

A strategy for Nature can be taken to be just one of the possible states. For example, in the problem of the fuses, the possible states for the purposes of the problem are just the possible numbers of defective fuses in the total lot of 1000. It is more usual to take a set of probability distributions over Z , the sample space, rather than the set of possible states as the space of Nature's strategies. It is assumed that to each one of the possible states of Nature there corresponds a unique probability distribution over Z , in that, given that Nature is in a particular state, then there is a definite, calculable probability that any one of the possible outcomes will be observed. For instance, if the number of defective fuses in the lot of 1000 fuses is n , then for all m , there is a definite probability that there will be m defectives among a sample of 20 drawn at random from the lot.

We have described strategies for the players, and the only factor still undefined is the payoff to the statistician resulting from the selection of a pair of strategies by the players. To give the definition, we must go into slightly more detail than we have so far done. This will involve distinguishing the cases in which the experiments themselves have no cost (where the performing of the experiment does not affect the total utility of the outcomes) from those in which the performance of the experiment must be reckoned into the utility of the final outcome. Since the operations in single-experiment decisions are all fixed in advance, these experiments may be assumed to have fixed costs; the cost may therefore be

neglected as a factor in the decision, and no generality is lost if we assume that these experiments have zero costs. The cost factor is not constant for sequential experiments, and since a wise selection of strategy by the statistician may be able to reduce it, a 'cost function' must be included as an explicit factor in studying sequential decision problems.

3.4.1 Payoff Function and Strategies for Single-Experiment Games

Let A be the set of terminal actions for the statistician, let Z be the sample space for the experiment he is to perform, let \mathcal{Y} be the set of states which Nature can be in, and let M be the payoff function such that $M(a, \mathcal{Y})$ is the payoff in utilities to the statistician resulting from action a by him and state \mathcal{Y} of Nature. A strategy (called a pure strategy to distinguish it from mixed strategies, which will be described later) for the statistician is a decision function, d , telling him what action to take for each of the possible outcomes of the experiment. Let D be the set of all such decision functions; D , then, is the space of pure strategies for the statistician.

Associated with each state, \mathcal{Y} , of Nature, there is a probability distribution $\omega_{\mathcal{Y}}$ over Z , such that if Nature is in state \mathcal{Y} , then $\omega_{\mathcal{Y}}(z)$ is the probability that z will be the actual outcome of the experiment. Let ' Ω ' be the set of all the probability distributions over Z defined in this way; then Ω is the space of strategies for Nature.

We must now define the payoff function corresponding to a choice of strategy d by the statistician and $\omega_{\mathcal{Y}}$ by Nature.¹ If Nature acts according to $\omega_{\mathcal{Y}}$ then for each outcome z , there is a definite probability $\omega_{\mathcal{Y}}(z)$

1. We shall assume that all of the distribution functions, $\omega_{\mathcal{Y}}$, corresponding to states of Nature are distinct.

that z will occur. The action taken in this event is $d(z)$, and the corresponding payoff $M(d(z), \varphi)$. Since each of these outcomes has a definite probability and is associated with the utility $M(d(z), \varphi)$, the compound Bernoullian utility due to strategies d and ω_φ is the sum of the payoffs for the particular outcomes, multiplied by the probabilities that those outcomes will be observed. Let $M(d, \omega_\varphi)$ denote the payoff due to strategies d and ω_φ , then

$$M(d, \omega_\varphi) = \sum_{z \in Z} M(d(z), \varphi) \omega_\varphi(z).$$

Besides the set of pure strategies, D , for the statistician, there is the corresponding set of mixed strategies, defined exactly as in 3.3. Having defined the payoff function now for pure strategies, its domain of definition is extended to cover mixed strategies also, as shown in 3.3. There is a second way of defining mixed strategies for single experiment games, which is perhaps only slightly more convenient to use than the first. This method consists first of extending the set of terminal actions, A , to include not only discrete acts, but probability combinations of those acts. If we were still working with a decision problem under uncertainty, in which the statistician's pure strategies are just the elements of the set A , then the extension we are considering would be just the set of mixed strategies over A , considered as a set of pure strategies. Let the extension of A be denoted A^* . As we have shown, the payoff function can be defined for these mixtures (we avoid referring to the members of A^* as mixed strategies, since we are reserving that term for mixed strategies over the domain of decision functions). We may now consider A^* as the set of terminal actions for the statistician, and consider decision function d^* which pick out for each possible outcome of the experiment a particular member of A^* to be taken if that outcome takes place. Formally, this function d^* is defined over the sample space, Z , and takes value in A^* , such that if the statistician

acts according to $d^*(z)$, and the outcome of the experiment is z , then he must take action $d^*(z)$, which is a risk combination of the basic set of action, A . We shall call strategies of this second kind mixed strategies also.

We may imagine the difference between the two kinds of mixed strategies as lying in the fact that in one case we define pure strategies as decision functions mapping the sample space, Z , into the set of 'pure' actions, A , then form the mixed strategies by taking probability combinations of the pure decision functions; in the second case, we first form all probability mixtures of the action space, A , then form mixed strategies as decision functions mapping Z into A^* , the set of probability mixtures of A . In either case, it is easy to show that the payoff function can be defined for that kind of strategy. It can be further shown that the two sets of mixed strategies are equivalent in the sense that for any mixed strategy d_1^* of the first kind, there exists a mixed strategy, d_2^* of the second, such that for all strategies, W of Nature, $M(d_1^*, W) = M(d_2^*, W)$, and vice versa. Let us denote the set of mixed strategies, whether of the first or second kind, D^* .

To illustrate the payoff functions and mixed strategies we return to the example of making a decision about the lot of 1000 fuses. A , the set of terminal actions, consists of the two alternatives 'accept' and 'reject'; let these be a_1 and a_2 respectively. The set of states of Nature are just the number of possible numbers of defective fuses in 1000; that is, Ψ is the set $\{0, 1, 2, \dots, 1000\}$. In order to formulate this problem, we must assume that a definite payoff $M(a_i, \varphi)$ corresponds to an action a_i and a state, φ . Suppose that the utility resulting if the lot is accepted is proportional to the number of non-defective fuses in it, and the expected

utility if it is rejected is the same as the utility of a lot containing 500 defectives. Then, we can assume M satisfies the equations:

$$M(a_1, n) = 1000 - n$$

$$M(a_2, n) = 500$$

where 'n' in the above equation denotes the state of Nature in which the lot contains n defective fuses.

The sample consisting of selecting 20 fuses at random from the lot and counting the number of defectives is a single experiment test, and we proceed to construct the set of strategies for the two players, and the corresponding payoff functions. The sample space for this experiment is the set of possible outcomes for the experiment, which is just the set of possible numbers of defective fuses in the sample of 20; this is the set $\{0, 1, \dots, 20\}$ which is denoted Z . A pure strategy for the statistician is a decision function which tells him for each possible outcome of the sample whether to accept or reject; it is a function, d , defined on Z and taking values in A , such that $d(m)$ is a_1 or a_2 according as this decision function directs him to accept or reject if he find m defectives in the sample. The set, D , of pure strategies for the statistician, is the collection of all such decision functions, and it is easy to show that there are just 2^{21} such strategies. The strategies for Nature are the probability distributions corresponding to the states of Nature, giving the likelihood that any particular outcome of the experiment will be observed if Nature is in the given state. Let W_n be the distribution over Z corresponding to the state in which the lot of 1000 contains n defective fuses; then $W_n(m)$ is the probability of finding m defective fuses in the sample of 20, given that

there are n defectives in the lot of 1000.¹ The set, Ω , consisting of all such distribution functions is the space of strategies for Nature. There are 1001 strategies, corresponding to the states $n = 0, 1, \dots, 1000$.

We can now use the formula of page 84 to calculate the payoff to the statistician resulting from the choice of strategy d by him, and ω_n by Nature. Since A has only two members, a_1 and a_2 , this formula reduces to:

$$M(d, \omega_n) = \left[\sum_{Med^{-1}(a_1)} \omega_n(m) \right] M(a_1, n) + \left[\sum_{Med^{-1}(a_2)} \omega_n(m) \right] M(a_2, n),$$

or, substituting the particular values given for the payoff function, we get

$$M(d, \omega_n) = (1000 - n) \cdot \sum_{Med^{-1}(a_1)} \omega_n(m) + 500 \cdot \sum_{Med^{-1}(a_2)} \omega_n(m).$$

It will be observed that the two sums on the right-hand side of the above expressions are just the respective probabilities the lot will be accepted or rejected, given that the statistician acts according to strategy d and Nature according to strategy ω_n .

The mixed strategies of the first kind are just the possible probability distributions over the set of pure strategies, P . If δ is such a mixed strategy, then the statistician is to use a random device in determining which of the pure strategies to use, such that it gives a probability $\delta(d)$ of choosing pure strategy d . Under these circumstances, the payoff function for mixed strategy δ and strategy ω for Nature is:

1. It is easy to show that for all $n = 0, 1, 2, \dots, 20$, $\omega_n(m)$ satisfies the equation

$$\omega_n(m) = \frac{980!}{1000!} \cdot \frac{n!}{(n-m)!} \cdot \frac{(1000-n)!}{(980-m-n)!} \cdot \frac{20!}{m!(20-m)!}.$$

$$M(s, \omega) = \sum_{d \in D} s(d) M(d, \omega)$$

as follows directly from the equation given in 3.3 defining the payoff function for mixed strategies.

To define the mixed strategies of the second kind, we must consider the probability combinations of the basic set of actions, A . Since A consists of just two actions, acceptance and rejection, the probability combinations of the actions can be represented by just two numbers: λ and $1-\lambda$, where λ is the probability of taking action a_1 (acceptance) and $1-\lambda$ is the probability of taking a_2 (rejection). Let A^* be the set of all these probability combinations of A . It is convenient for this example to represent each member of A^* as a single number λ , where if λ represents the compound action, a^* in A^* , then a^* consists of taking a chance λ of performing a_1 , and $1-\lambda$ of performing a_2 . A mixed strategy of the second kind is a decision rule which picks out a certain member of A^* to be performed for each outcome of the experiment. If the statistician follows mixed strategy d^* , and if the outcome of the experiment is z , then he must take the action corresponding to $d^*(z)$; i.e., take action a_1 with probability $d^*(z)$ and a_2 with probability $1-d^*(z)$. If the statistician has determined on a mixed strategy d^* to follow, and Nature follows strategy ω_n , then there is a definite probability that action a_1 will be taken, which is in fact the sum:

$$\sum_{z \in Z} \omega_n(z) d^*(z)$$

and a corresponding probability that a_2 will be taken. The payoff to the statistician corresponding to the strategies d^* and ω_n satisfies the equation:

$$M(d^*, \omega_n) = \left(\sum_{z \in Z} \omega_n(z) d^*(z) \right) M(a_1, n) + \left(1 - \sum_{z \in Z} \omega_n(z) d^*(z) \right) M(a_2, n)$$

which reduces to

$$H(d^* | W_n) = \left(\sum_{m=1}^n W_n(m) d^*(m) \right) (1/100 - n) + \left(1 - \sum_{m=1}^n W_n(m) d^*(m) \right) .500.$$

3.4.2 Sequential Games

The basic experiments of sequential games are like those for the single-experiment games previously described; however, in the sequential game, the experiment is analyzed into a sequence of sub-experiments, and the statistician is allowed to terminate experiments at any point in the sequence if he so desires, leaving the rest of the sub-experiments unperformed.¹ Thus, the experiment described in the last section can be analyzed into a sequence of twenty sub-experiments, each consisting of selecting one fuse, testing it, and noting whether or not it is defective. To transform this experiment into a sequential game, it is only necessary to allow the statistician to stop at any point in the sequence and make his final decision at that time. There would be no point in stopping the tests before the end if the test themselves cost nothing, and therefore sequential experiments are of practical interest where the tests have some positive cost (or negative utility). If, for example, it were necessary to destroy a fuse in order to test it, the statistician would have a practical interest in reducing the number of tests as much as possible.

The possibility that the statistician may terminate the tests at any point in the sequence (or even before the sequence has begun) greatly enlarges the range of strategies available to him. As we have seen, a strategy in a single experiment game is a decision function directing what action should be taken in the event any outcome of the experiment is observed. In a sequential game, a strategy must include a rule which tells the statistician whether to continue experimenting or quit at some point in a sequence

of observations, and it must include a rule which tells the statistician what action to take for each possible way the experiments may terminate.

Let us now try to formalize the notion of a strategy for the statistician. As before, let A be the set of final action. Let Z be the set of outcomes of an overall compound experiment, which is analyzed into a sequence of sub-experiments or "component" experiments $1, 2, \dots, k$, whose outcomes comp as the sets Z_1, Z_2, \dots, Z_k . Each outcome, z in Z , of the compound experiment is a sequence of outcomes, z_1, \dots, z_k , of the component experiments. In general, we shall write

$$z = // z_1, \dots, z_k //$$

where z is the compound outcome which corresponds to the component outcomes z_1, \dots, z_k .¹ In what follows it will be convenient to adopt the following viewpoint with respect to the compound experiments whose outcomes are z .

Before any of the sub-experiments are begun, it is assumed that the compound experiment has a predetermined outcome, $z = // z_1, \dots, z_k //$. Performance of the sub-experiments reveals in order what the components of z are.

Cessation of the sub-experiments before the final one is performed means that the remaining components of z remain unknown to the statistician, but these outcomes are assumed to exist nevertheless. The adoption of this convention, although it may be offensive logically, greatly simplifies the statements of the definition of strategy for the statistician. The essential point is that definitions must not depend on unobserved components of the outcomes.

Part of the statistician's strategy consists of a rule telling him whether to stop or continue at any point in his experimentation. Such a rule is called a sampling plan. Formally this will be represented by a function, g , of two variables, such that if the outcome of the compound

experiment is s , then the statistician is to continue after observing the first i components if $g(i, s) = 0$, and stop if $g(i, s) = 1$. Thus g is a function which tells the statistician for each compound outcome, s , how many components he should observe: namely, he is to continue observing components x_1, x_2, \dots until he arrives at some i for which $g(i, s) = 1$; he is to observe x_i and then stop. We shall place two restrictions on the function g . The first is the obvious one that for fixed i , $g(i, s)$ must not depend on the components of s after x_i : this for the reason that $g(i, s)$ is supposed to direct the statistician whether to stop or continue, and this cannot be permitted to depend on components about which the statistician is still in ignorance. Secondly, we stipulate that for all s , $g(i, s)$ is 1, (i.e., directs a stop) for exactly one $i = 0, 1, \dots, k$. With this restriction, $g(i, s)$ has an interpretation which will prove convenient later: namely, $g(i, s)$ is the conditional probability that, given that the outcome of the compound experiment is s , the number of components observed will be i ; that is, that the statistician will observe x_1, \dots, x_i , then stop. Let the set of sample plans be \mathcal{G} .

The way in which a sampling plan function, g , would be used in practice would be to begin the sequence of experiments, and note their outcomes, x_1, x_2, \dots until a point was reached for which the outcomes observed were in fact the first i components of some compound outcome s for which $g(i, s) = 1$. Let us return to the problem of the fuses for a concrete example. A sampling plan is represented by the instructions: "test fuses one at a time, noting them as non-defective or defective, and stop after either eleven defectives are found or twenty fuses have been tested." There are twenty sub-experiments in this case, each one of which has two possible outcomes: non-defective or defective, which we denote 0 and 1 respectively.

Hence a compound outcome is a matrix of twenty components, which are 0's and 1's. The function g in this case satisfies the conditions:

$$g(i, z) = \begin{cases} 1, & i < 20 \text{ and } z_1 + \dots + z_i = 11 \text{ and } z_i = 1, \\ 1, & i = 20 \text{ and } z_1 + \dots + z_i \leq 11 \\ 0, & \text{otherwise.} \end{cases}$$

The reader can easily verify: (i) that the function defined above satisfies the two restrictions stated on page 16 and (ii) that it directs continuation or stopping in exactly the same situations that the verbal instructions did.

The second part of the statistician's strategy consists of a rule which prescribes what action he shall take for each of the possible ways the sequence of experiments may terminate. We shall call such a rule a decision rule. We imagine that the compound experiment may have any outcome, z in Z , and that the sampling plan may demand that the sequence of observations may be terminated at any component, z_i , $i = 0, 1, \dots, k$. A decision rule, then, must tell the statistician for any possible z , and any possible stopping point, z_i , what action to take.¹ Formally a decision rule is represented by a decision function, d , of two variables, such that for all z and for $i = 0, 1, \dots, k$, $d(i, z)$ is the action (i.e., $d(i, z)$ is in A) which the statistician is to take if the outcome of the compound experiment is z , and the sampling plan requires that the sub-experiments stop after the first i components have been observed. We must place a restriction on the allowable decision functions analogous to that we placed on the sampling plan functions, g . That is, for fixed i , $d(i, z)$ must not depend on the unobserved components z_{i+1}, \dots, z_k . Let the set of decision functions be denoted D .

1. In most cases, a decision rule of this type will prescribe actions for sequences of observations which cannot arise, because of the fact that the sampling plan chosen requires that either the observation be stopped before that point is reached or continued past that point.

A strategy for the statistician is a sampling and a decision rule, which are formally represented by a pair of function g in G and d in D . If the statistician acts according to g and d , and the outcome of the compound experiment is s , then he is to observe components s_1, s_2, \dots until some sub-experiment s_i is reached for which $g(i, s) = 1$, then to stop and take action $d(i, s)$.

In practice, if the statistician acts according to a strategy represented by g and d , he observes outcomes of the sub-experiments,¹ noting them as s_1, s_2, \dots etc., until some sequence s_1, \dots, s_i has been observed which contains the first i components of a compound outcome s for which $g(i, s) = 1$, then experimentation is to stop and the action to be taken is $d(i, s)$. A strategy is represented by the instructions: "test fuses drawn one by one at random, until either eleven defectives have been found or twenty fuses have been tested, then accept the lot if less than eleven defectives have been found and reject it otherwise." We have already constructed the function g representing the sampling plan for this example. The set of final actions is just acceptance and rejection, which we denote a_1 and a_2 respectively; then for all i and s , $d(i, s)$ must be either a_1 or a_2 . The reader may easily verify that the function d_1 defined

$$d(i, s) = \begin{cases} a_1 & \text{if } s_1 + s_2 + \dots + s_i < 11 \\ a_2 & \text{if } s_1 + s_2 + \dots + s_i \geq 11 \end{cases}$$

is actually the desired decision function. We note, too, that this definition refers only to the first i components of s , as required.

1. It is understood that if $g(0, s) = 1$, then the sequence is stopped before it has been begun.

We have as yet said nothing about the space of strategies for Nature. The set of states of Nature is \mathcal{V} , and, as we showed in 3.4.1, to each φ in \mathcal{V} corresponds a probability distribution W_φ over \mathcal{Z} , such that $W_\varphi(z)$ is the probability that the outcome of the compound experiment is z , given that Nature is in state φ . It is convenient to generalise this and consider functions W_φ of two variables such that $W_\varphi(i, z)$ is the probability that the outcomes of the first i sub-experiments are $z_1 z_2 \dots z_i$ given that Nature is in state φ . If i is held fixed, then $W_\varphi(i, z)$ actually is a probability distribution giving the probabilities for sequences of outcomes of the first i sub-experiments. The set of these functions for φ in \mathcal{V} is the space of Nature's strategies, and is denoted Ω .

All that remains now is to define the payoff corresponding to a strategy represented by g and d for the statistician and W_φ for Nature. As in the case of the single experiment game, we assume that there is an underlying payoff function M , such that for each action a , and state φ $M(a, \varphi)$ represents a payoff in utilities to the statistician. Besides this, however, there is a cost, which must be subtracted from the final payment as the price of experimentation. We assume that the cost is represented by cost function, c , of two variables, such that $c(i, z)$ represents the loss to the statistician in utilities if the outcome of the compound experiment is z , and he performs the first i sub-experiments. As before, we require that for fixed i , $c(i, z)$ is independent of the components $z_1 + z_2 + \dots + z_i$. We assume furthermore that performing each additional sub-experiment adds to the cost, or at least does not decrease it, and that performing no experiment cost nothing. Formally,

$$c(0, z) = 0,$$

$$\text{if } i > j, \text{ then } c(i, z) \geq c(j, z).$$

The payoff is now simply defined. If the statistician chooses strategy g , d , and Nature chooses ω_p then the probability that the sequence of sub-experiments actually observed is z_1, z_2, \dots, z_k is

$$g(i, z) \omega_p(i, z) \quad (1)$$

and the payoff associated with this sequence of observations and pair of strategies is

$$M(d(i, z), \psi) - c(i, z).$$

The payoff function, $\rho(d, g, \omega_p, c)$, due to strategies g, d , and cost function c is defined by the equation:

$$\rho(g, d, \omega_p, c) = \sum_{z \in Z} \frac{1}{L} \sum_{i=0}^k M(d(i, z), \psi) - c(i, z) (g(i, z) \omega_p(i, z)).$$

1. Recall that $g(i, z)$ can be interpreted as the conditional probability that z_1, z_2, \dots, z_k will be observed given that the actual sequence of outcomes is z_1, z_2, \dots, z_k , and that the Statistician acts according to g .

3.5 Decision Principles

So far we have set forth the formalism for representing a wide variety of problems without actually giving the solutions to any of these problems. A "solution," in the sense that we are using that term here, is a rule which tells player 1, in whose decision we are interested, which strategy to choose from among those available to him.¹ A rule which directs what decision to make in a large class of games is called a decision principle. Decision principles must be distinguished from decision functions in statistical games. A decision function is a strategy in a particular statistical game, whereas a decision rule is a rule which picks a strategy in a large class of games. In this section, we shall discuss some of the considerations which lead to the selection of some rational decision principles which have been applied to certain classes of problems, and the rationales behind them. Our discussion will be very limited for two reasons: first, because the general theory of decisions, including statistical decisions, is very complex and would take us far from our main objective;² second, our main objective is to show what role utility theory plays in the general theory of decisions, and, in one sense, this has already been accomplished by showing how the payoff functions are defined in terms of the Bernoullian utility functions of the players. The only way that utilities enter into the formal decision problem is in their effect on the payoff functions, and once the payoff functions are defined over the strategy spaces of the players, the decision problem is defined. However, discussion of the decision principles throws light on utility theory itself, since, as we shall see, the

1. The definition of 'solution' in von Neumann's theory of n-person games differs from our usage here. The von Neumann solution of a game describes a certain property which it can be argued the game must have if all the players play rationally, but does not describe what the actual strategies of the players will be, and hence it does not provide a decision principle for the players.

2. The reader is once again referred to the Theory of Games and Statistical Decisions by Garahick and Blackwell for an introduction to the technical side of this subject. [6]

choice of a decision rule which provides the solution to a decision problem should be based on arguments analogous to those advanced in justification of the axioms of Bernoullian utility. That this should be the case is obvious if we recall the Bernoullian utility itself is a theory of rational decisions, and is meant to provide a decision procedure in situations in which the outcome of the 'game' depends only on the action taken by the person making the decision, and on chance factors of which the probabilities are known. Thus, examination of the arguments supporting the choice of a decision principle will serve to clarify the conceptual basis of utility theory itself, and will help to expose some of the restrictions which must be placed on the applications of Bernoullian utilities.

Utility theory itself provides us with the first example of a decision principle, namely: in choices among alternatives involving only risk factors, choose that action with the highest utility. The rationale for this principle is given in the arguments which justify the axioms for Bernoullian utility. The 'games' for which Bernoullian utility provides the solution are essentially all the 'one-person' games (see page 72). Once past the one-person games, we shall find that, except for a very small class of two-person games, there are no decision principles which provide solutions as satisfactory as those Bernoullian utility provides for one-person games. Beyond the one-person games, the field of decision problems may be conveniently subdivided as follows: (1) two-person games in which both players are rational (i.e., play for self-interest); (2) n -person games in which all the players are rational; (3) two-person games involving one rational player against a non-rational opponent (Nature); (4) n -person

games involving both rational and non-rational players.^{1,2} It is only for a certain sub-class of (1) that there exists a decision principle which is generally accepted as providing satisfactory solutions. This is the two-person game called zero-sum in which the interests of the two players are diametrically opposed. We shall examine this type of game and the corresponding decision principle below. For the non-zero-sum games of class (1) there is at present no satisfactory theory, although some starts have been made. It is possible to define the notion of 'zero-sum' for games of class (2) analogously to its definition for the games of class (1), and there is a theory for the zero-sum games of class (2). However, this theory suffers from two defects: first, it is based on the assumption that the players in the game will form themselves into two coalitions which then play as if they were playing a two-person game; and second, this theory does not "solve" the decision problem in the sense that we are using that term. Class (3) includes those statistical games whose strategies were described in 3.4. There is an extensive theory of statistical decision procedures, but there is no decision principle which is generally recognized as providing a satisfactory solution to all problems of this class. Class (4) is mentioned only for the sake of completeness; so far as is known to the author, there is at present no theory for games of this class.

3.5.1 Two-Person Games Between Two Rational Players: the Minimax Principle

The most obvious games of the class we are considering are two-person parlor games like checkers and chess, and two-person forms of games

1. We have shown previously (p. 74) that if one or more of the non-rational players plays according to random strategies whose probability distributions are known to the players, then the game can be reformulated without these players. Hence it is sufficient to consider games in which the non-rational players represent uncertain factors.

2. 'N-person' game here refers to games with more than two players.

like poker and craps. Less obvious but still with enough gamelike features to make a game-theoretic approach seem fruitful are certain other situations involving essentially two competitors, such as duopoly and duopsony, military problems, bargaining problems, and some kinds of duels. The general decision problem is to find a rule which prescribes what strategy player 1 (and by analogy, player 2) should use in all games of the class under consideration. Before stating any specific proposals for decision principles, we shall discuss some general considerations involved in the choice of one.

The main feature distinguishing games in which one or more of player 1's opponents are rational from those in which none is, is that the rational opponents will attempt to anticipate player 1's strategy in order to profit by this anticipation. Player 1 will, of course, try to anticipate the opponents' choices also, and in doing this he must take into account their estimates of him. This reasoning may seem to complicate the game between rational players to the point where the problem becomes unmanageable; however, there is one simplifying assumption we can make which affords us a certain amount of guidance in seeking a rational decision principle. We may assume that if there is a rational way for player 1 to play the game, it is the rational decision procedure for his opponents as well. If there is only one rational way to play then, for a given game, each player will know what strategy the opponent will use, and choose his to make the most of it. We may then require that even if player 1 knows that his opponents will play according to the same decision rule that he uses, he will have no reason to change his own strategy. To formalize this assumption, let S_1 and S_2 be the strategy spaces for players 1 and 2, and let M_1 and M_2 be their respective payoff functions. Suppose that under a contemplated decision

rule, player 1 should choose s_1 as his strategy and player 2, using the same rule, should choose s_2 . The payoff to player 1 under these circumstances is $M_1(s_1, s_2)$. We should like to require that, even if player 1 should know in advance that player 2 will choose s_2 , he will still have no reason not to choose s_1 . However, if there is some strategy, t_1 , for player 1 such that

$$M_1(t_1, s_2) > M_1(s_1, s_2)$$

then, player 1 should clearly prefer t_1 to s_1 , if he knows player 2 will choose s_2 . Therefore, we should expect our decision procedure to be such that for all strategies, t_1 , for player 1,

$$M_1(t_1, s_2) \leq M_1(s_1, s_2)$$

and by analogous reasoning, for all strategies, t_2 for player 2,

$$M_2(s_1, t_2) \leq M_2(s_1, s_2).$$

A pair of strategies, s_1 and s_2 , satisfying the two foregoing conditions, are called equilibrium strategies. Each of a pair of equilibrium strategies has the property that it is the best strategy to use if the opponent chooses the other. We should like our decision procedure to have the property that if both players follow it, they will always come up with pairs of equilibrium strategies.

In general, we may say that the condition that a proposed decision procedure always give equilibrium strategies in the class of games to which it is applied is a necessary condition for its adequacy, but is not sufficient. The class of two-person games between rational players may be broken up into three sub-classes as follows: games in which there is no pair of equilibrium strategies; games in which there is exactly one pair of equilibrium

strategies; and games in which there are two or more pairs of equilibrium strategies. For games of the first kind, it seems there is no rational solution to the decision problem as it has been stated, although we shall find that if the class of admissible decision procedures is enlarged, some of these games with no equilibrium strategies will prove to have solutions. As an example of a game with no equilibrium strategies, consider the game in which the two players have two strategies each, s and t for player 1 and x and y for player 2, and their payoff functions are shown in the following table:

		player 2's strategies	
		x	y
player 1's strategies	s	0/2	3/1
	t	1/0	2/1

Two numbers are given at each place in this table, the first indicating the payoff to player 1 and the second being the payoff to player 2 for the corresponding strategies by the two players. For example, $M_1(s, x) = 0$ and $M_2(s, x) = 2$, since 0/2 is the entry in the table for strategy s by player 1 and x by player 2.

The reader can easily convince himself by examining the table that there is no pair of strategies for the players such that each is the best to use against the other strategy of the pair. In this game it appears that there is no single decision rule which is rational for both players to follow, for if there were, both players would know it and thus be enabled to anticipate the other's strategy in order to choose their own best strategy.

In case there is just one equilibrium pair in the game, these are the only strategies which are eligible to be considered as rational solutions

to the decision problem. However, even then, the equilibrium pair may not be intuitively acceptable as a solution. Consider the game in which player 1 has two strategies, s and t , and player 2 has three strategies, x , y , and z , and the payoffs are represented in the following table:

		player 2's strategies		
		x	y	z
player 1's strategies	s	0/0	3/-2	0/-1
	t	-1/-1	2/1	1/0

In this game the strategies s and x are the only equilibrium strategies, as the reader can verify by checking each of six possible pairs in the table. However, it would seem that for many reasons the pair t and y would be preferable, since both players actually receive more from these two than they do from the pair s, x . However, y is not a "safe" strategy for player 2, since player 1 would have an incentive not to choose t if he knew player 2 would choose y . The equilibrium strategies are the only "safe" strategies, since the players know that the opponent, even if he knows what the first player's strategy will be, has no incentive to change. However, whether we wish to accept the equilibrium strategies as solutions in games with only one pair of them is a matter which may be questioned. If we choose to reject the equilibrium strategies as solutions, then we must say, as with the games with no equilibrium strategies, that the decision problem has no solution for these games.

For games in which there is more than one pair of equilibrium strategies, the decision problem is even more confused. For such games the

players cannot be sure that the strategy they pick is a member of the same equilibrium pair that the opponent's strategy is.

Consider the following game:

		player 2's strategies	
		x	y
player 1's strategies	s	1/1	-10/-10
	t	-10/-10	0/2

here the pairs s, x and t, y are both equilibrium pairs. It would seem that player 1 should prefer to choose the pair s, x and player 2 should prefer t, y since player 1 gets more in the first and player 2 gets more in the second. However, if player 1 follows his inclination and chooses s , and player 2 chooses y , they will both end up with -10, which is much worse for both than either of the equilibrium pairs.

It is possible to treat some games in which there is more than one equilibrium pair as if they had only one equilibrium pair. These are games in which if s_1, s_2 and t_1, t_2 are equilibrium pairs, then s_1, t_2 and s_2, t_1 are also equilibrium pairs. For these games it is easy to show that the payoffs from all the equilibrium pairs are the same:

$$M_1(s_1, s_2) = M_1(s_1, t_2) = M_1(t_1, s_2) = M_1(t_1, t_2),$$

and

$$M_2(s_1, s_2) = M_2(s_1, t_2) = M_2(t_1, s_2) = M_2(t_1, t_2)$$

In this case, it makes no difference to the outcome which of the possible first members of the equilibrium pairs player 1 chooses, and which of the second members player 2 chooses, since this pair of chosen strategies must also be an equilibrium pair, and the payoffs from this pair are the same

as from any other equilibrium pair.

We have stated the decision problem as asking for a rule which tells player 1 which strategy to choose in all the games of a certain class. We have found that for many of the two-person games which we have been considering there appears to be no rational solution to the decision problem either because there is no pair of equilibrium strategies, or because there are too many such pairs. It is possible to rephrase the decision problem somewhat so as to enlarge the class of admissible decision procedures. Rather than asking for a decision rule which tells player 1 unequivocally which strategy to choose in any given game we may ask instead for a decision rule which tells player 1 to follow some procedure which will in turn tell him what strategy to choose. We shall not consider all such procedures, but confine our attention to a special type called "randomized" procedures, or mixed strategies. A mixed strategy may be regarded as a procedure, involving the use of random devices, whose outcome tells player 1 which pure strategy to use. Hence, a decision rule which directs player 1 to follow a certain mixed strategy in a game does not tell him unequivocally which pure strategy to follow, but directs him to use a certain random procedure, the outcome of which does tell him which pure strategy to use.

Of course, our concentration on mixed strategies is not accidental. We have seen that the payoff functions may be extended to include mixed strategies, hence that the mixed strategies may be regarded in turn as pure strategies in a game whose payoffs are given by the extended payoff functions. Therefore, all the considerations relating to equilibrium strategies and rational decision principles apply directly to these extended games. The important point is that some games which have no equilibrium pairs in the space of pure strategies have such pairs in the extended game of mixed

strategies. In these games, the mixed strategies constituting the equilibrium pair in the extended game may be considered as eligible candidates for solutions to the decision problem for the original game.

There is one particularly important class of games for which the mixed strategies always contain equilibrium pairs, and if they contain more than one such pair, these pairs always satisfy the condition of equivalence stated on page 11. These are the so-called zero-sum games. Zero-sum games are those in which the interests of the players are diametrically opposed in the sense that what benefits one must hurt the other. This condition of strict opposition can be stated as follows: if s_1 and t_1 are two strategies for player 1, and s_2 and t_2 are two strategies for player 2, then $M_1(s_1, t_1) \geq M_1(s_2, t_2)$ if and only, if $M_2(s_1, t_1) \leq M_2(s_2, t_2)$. If the above condition is satisfied where s_1, s_2, t_1 and t_2 are arbitrary mixed strategies for a given game, then it is easy to show that it is possible to choose eligible utility functions for the players and associated payoff functions, M_1^0 and M_2^0 such that for all strategies s for player 1 and t for player 2,

$$M_1^0(s, t) + M_2^0(s, t) = 0,$$

hence the term "zero-sum" for these games. It is worthwhile to note that a game may satisfy the condition of strict opposition for its pure-strategy spaces, but not for its mixed-strategy spaces. The condition of strict opposition implies that all pairs of equilibrium strategies of a game are equivalent in the sense of page 104, however, it does not by itself guarantee that an equilibrium pair exists. If the strict opposition carries over to the mixed strategy spaces of a game (as well) - in other words, if the game is zero-sum - then there exists at least one pair of equilibrium strategies in the spaces of mixed strategies. Equilibrium pairs of mixed strategies in zero-sum games are called minimax strategies.

There is very good reason to accept the minimax strategies in zero-sum games as rational solutions to the decision problem for these games. As we have noted, a pair of minimax strategies for players 1 and 2 have the property of being equilibrium strategies; i.e. each is best to use against the other, so that a player will have no reason to change his strategy even if he finds out what strategy the opponent is using, or the opponent finds out what strategy he is using, or both find out what the other is using. Furthermore, there is no possibility in the zero-sum game that there are other pairs of strategies which give both players more, as in the example on page 9, since any shift which benefits one player must hurt the other. Finally, in certain classes of games, in which the rational mode of play seems very clear (such as checkers, chess, tic-tac-toe), the minimax strategy coincides precisely with these rational modes of play.

The minimax principle (i.e. the principle that players should play according to the minimax strategies) furnishes an intuitively acceptable solution to the decision problem for zero-sum games. It can be extended to cover the slightly larger class of games which satisfy the condition of strict opposition over the space of pure-strategies, even though not over the space of mixed strategies. In this case, if the game has an equilibrium pair, it is equivalent to all other equilibrium pairs, and furnishes an intuitively acceptable solution. Except for these relatively restricted classes of games, however, there are no generally accepted decision principles.

We may mention here two theories which are meant to deal with some of the non zero-sum games. One is Nash's theory of bargaining problems,¹ which can be regarded as two-person games, and the other is Raiffa's theory

1. Nash, John F. [17]

of arbitration procedures,¹ which can be regarded as procedures, like mixed strategies, which enable players to arrive at a decision as to which pure strategies to choose.

Before leaving the two-person game between rational players we return to a point raised previously (pages 54 and 79) about the impossibility of treating the strategy spaces themselves as outcomes over which Bernoullian Utilities can be defined. It is clear that Bernoullian Utilities cannot reflect rational preferences for pure and mixed strategies in zero-sum games, for it often happens that it is rational to choose a mixed strategy, but not rational to choose any of the pure strategies of which it is a randomization. This violates axiom A.4, (page 49), hence indicates that the justification given for axiom A.4 (see page 50) is not valid in this case. A.4 states that if x and y are two outcomes such that x preferred to y , and $\langle \alpha x, (1-\alpha) y \rangle$ is a random combination of x and y , then x is preferred to $\langle \alpha x, (1-\alpha) y \rangle$ and $\langle \alpha x, (1-\alpha) y \rangle$ is preferred to y . The justification of the first part of this is that the final outcome of $\langle \alpha x, (1-\alpha) y \rangle$ is either x or y , and since x is preferred - or - indifferent to x , and is preferred to y , then it should be preferred to the random combination. This justification rests on the still more fundamental assumption that the actual act of randomization does not affect the outcome which finally results: i.e., that it makes no difference to the person whose preferences we are considering whether he simply receives outcome x directly, or as a result of taking the risk combination $\langle \alpha x, (1-\alpha) y \rangle$. However, it does make a difference in a game whether player 1 uses strategy s outright or as a result of following some mixed strategy, say $\langle \alpha s, (1-\alpha) t \rangle$. One indication of this difference lies in the difference in the concept of equilibrium strategy as applied to pure and mixed strategies. A pair of strategies is in equilibrium if neither player would have any incentive to change if he knew the other's strategy

1. Raiffa, Howard [22]

or the other knew his. If the players are using mixed strategies, it is impossible for either player to know in advance what his or his opponent's pure strategy will be. Hence the definition of equilibrium strategies is different for mixed strategies. It would be extremely difficult to trace through the exact consequence of this change in meaning to show why there are equilibrium mixed strategies in games in which there are no equilibrium pure strategies, and why a mixed strategy may be preferred to all the pure strategies of which it is compounded, and we will not attempt it here. It is sufficient to point out the significance of randomization for the equilibrium concept, and the consequence that randomization of itself affects the situation in which it is applied.

3.5.2 Strategies in n-person games, all players rational

The notion of equilibrium strategies carries over very naturally to the n-person game: strategies s_1, \dots, s_n for players $1, \dots, n$ are equilibrium strategies if each strategy, s_i , $i = 1, \dots, n$ is the best strategy for player i to use, given that the remainder of the players will all choose the other strategies of the set. In theory, any rational solution to the decision problem should be one such that the strategies chosen by the players of any particular game constitute an equilibrium set. Actually, for almost all n-person games, there are many sets of equilibrium strategies, and these are not all equivalent; hence the theory meets the same difficulties encountered in the two-person game with equilibrium pairs which are not equivalent.

Von Neumann has developed a theory of n-person games which depends on reducing n-person games to two-person games by assuming that the players form themselves into coalitions which then become 'super-players' in a two-person game. We shall not consider this theory in detail,¹ but shall comment

¹ See Luce, R. D., Survey of Games, Part III - n-Person Games, 7:8 No.5
Bureau of Applied Social Research, May 1954. [37]

an assumption made in the theory which has interesting implications in the foundations of utility. The games for which the von Neumann theory is designed are all assumed to be zero-sum in the general sense, that is, it is assumed that it is possible to pick eligible utility functions for the players, with corresponding payoff functions, U_1, \dots, U_n , such that for any choice of strategies, s_1, s_2, \dots, s_n ,

$$\sum_{i=1}^n U_i(s_1, \dots, s_n) = 0.$$

A second assumption on which the theory depends is that in some way it is possible for players within a coalition to make "transfers" of utility among one another so that the payoff to the coalition resulting from a given choice of coalition strategy may be distributed in an arbitrary way among the members. The foregoing assumption is often labeled the 'assumption of transferability' of utility, and is a basis for many attacks on the empirical applicability of von Neumann's n-person game theory. Stated as an assumption of 'transferability', of course, this assumption is false on logical grounds alone, since the word 'transfer' applies to physical objects, not to the numbers which are the values of the utility function. However, the assumption can be restated in empirically meaningful terms in such a way as to meet the requirements of von Neumann's theory. What is necessary is that there be acts which players can perform which result in utility changes to the players but for which the sums of their utilities before and after the act are the same. To take a concrete example, the act in question may be for the first player to hand the second a dollar bill. If, in the scale of utilities in which the game's payoffs are being computed, the change of utility to the player who receives the dollar is the negative of the change of utility to the player who gives the dollar (in other words, the sums of

their utilities before and after are the same), then the act of handing the dollar bill performs the type of function which is required by the von Neumann theory.

The von Neumann theory does not throw much light on the decision problem as we have stated it. It does, however, imply two principles of behavior (either descriptive or rational, depending on the basic interpretation of the theory). They state which coalitions can form, and how the payments will finally be distributed among the members of the coalition. These, of course, rest on the two assumptions mentioned above, and on the assumption that the utility 'transfers' demanded by the theory will actually be made. In a wide sense, all these assumptions come under the heading of decision theory, just as all voluntary behavior falls into this category; however, the discussion of these assumptions is too large a topic, and would take us too far afield to be included here.

3.5.3 Statistical Games

We shall not attempt in this section to discuss all or even a large percentage of the various decision principles which have been advanced for statistical games. In our previous discussions, the choice of a decision rule depended on some assumption about what strategy the opposing player would follow. However, in the statistical game, as in the ordinary decision under uncertainty, we expressly assume that the person making the decision has nothing to guide him in guessing what strategy his opponent (nature) will follow. The statistician may choose a final action after gathering statistical information as to the state of nature, such since his strategy is a statistical decision procedure,¹ he must choose this before he ever gathers his information. Therefore the statistical decision problem may be regarded as a special

1. See Sect. 3.4.

case of the decision problem under uncertainty, and we shall not expect that considerations as to the actual state of nature will play any part in the choice of a decision procedure for statistical games.

Many decision principles are used in practice; there is, however, one condition to which it is natural to require that all of them conform. This is the condition that they always select admissible strategies in the games to which they are applied. An admissible strategy, s , is one such that there is no other strategy, say t , which gives player 1 a payoff as high or higher no matter what strategy his opponent picks. An inadmissible strategy then, is a strategy for player 1 such that there is another strategy for player 1 which gives him a better payoff (or at least as good a payoff) no matter what strategy his opponent chooses. It is intuitively clear that player 1 should never play according to an inadmissible strategy, and hence that he should consider only those decision procedures which constitute admissible strategies.

One decision principle that is commonly used is a Bayes Principle. Suppose that D is the class of decision functions (i.e. strategies for player 1) for a given statistical game, Z is the sample space for this game, Ω is the set of probability distributions over Z corresponding to the possible states of Nature (i.e. Ω is the strategy space for Nature), and M is the payoff function for player 1. Player 1 may assume that Nature picks a strategy according to some random plan or probability distribution, ξ , and under these circumstances it is possible to define a utility for each of the decision functions in D , and solve the decision problem by choosing that function with the highest utility.¹ The utility of a decision function

1. See page 76.

function, d , if Nature uses random strategy ξ is

$$\sum_{\omega \in \Omega} M(d, \omega) \xi(\omega),$$

and it is only necessary to choose d so that the above expression is maximized. Such a decision procedure is called a Bayes procedure. It can be shown that any strategy selected by a Bayes principle is admissible.

It should be noted that there is no one Bayes procedure, since the distribution function, ξ , will be chosen, presumably, according to considerations relating to the particular game in question. Indeed, the decision problem for statistical games could be rephrased to ask what is a rational assumption to make about the distribution function ξ .

A decision principle which does "solve" the decision problem without references to any arbitrary factors such as the distribution function assumed by the Bayes, is the Minimax principle.¹ Formally stated, the Minimax principle says to choose d such that

$$\min_{\omega \in \Omega} M(d, \omega)$$

is a Maximum; i.e. to choose d so that the worst possible outcome to player 1 from any choice, ω , by Nature is a maximum. This is actually the same principle as the Minimax principle for the zero-sum two-person games, the difference being that as applied to statistical games, it does not have the same justification as it does for the zero-sum games. In the zero-sum games between two rational players there is good reason to believe that the opponent will choose his strategy so as to hurt his opponent the most, but there is no reason to believe that Nature will act in this way. The Minimax principle may be called conservative, since it picks a strategy which minimizes the possible loss to player 1. At the other end of the scale, it would be

1. This may be regarded as a special case of the Minimax solution defined for zero-sum two-person games.

possible to define a principle which always picked a strategy which maximized the possible gain to player 1. In between the extremely optimistic and extremely pessimistic principles are a variety of others such as the Minimax Loss principle. This principle directs that d be chosen so that the maximum loss (difference of actual payoff and best possible payoff) be a minimum.

All the above principles can be shown to pick admissible strategies, and there seems to be little reason to choose one in preference to the others. All these principles have been applied in practice and it is probably fair to say that which one is applied in a particular instance is a matter of the statistician's taste.

4. Descriptive Applications of Bernoullian Utility Theory

In section 3, we considered applications of utility theory to the decision problem in general. In our discussion there, we assumed that utility theory, and the various decision theories based on it, were theories of rationality - that they provided principles which were in some sense the best to follow in gaining the objectives whose values are given by the utility function. That discussion, including the different definitions of decision principles, can be carried over and applied to utility decision theory as theories of actual behavior. In fact, the descriptive applications of Bernoullian utilities are to just those areas which correspond to the theories of rationality described in section 3; i.e. to decision-making behavior. Unfortunately, the assumption the players, or the decision-maker, know what all the possible strategies are, and what the corresponding payoff functions are, and are able, in the case of game theory, to calculate the minimax solution, is all but fatal to any descriptive interpretation of these decision theories in situations of even moderate complexity. Therefore, we shall find that all the empirical applications of utility theory have been made in extremely simple (sometimes in artificially schematized) situations.

So far, the main empirical applications have been made to situations in which the actual payments at the outcome were in money, and where, as a consequence, the only utilities involved are for amounts of money. In the next two subsections we discuss two such applications.

We shall see that the attempt to apply Bernoullian utilities predictively brings up problems for which there is no counterpart in the interpretation of utility as a theory of rationality. To mention one of these problems, it is necessary to assume, if utility theory is to be used predictively,

that utilities remain constant over time, or that if they do not remain constant, then it is necessary to know the laws governing the way they change. In our discussion of the basic interpretation of utilities (Sec. 2.2), no such assumption was made, and in fact, we have noted some reasons why utilities should not be applied to sequences of decisions. We shall discuss some of these problems in sections 4.3.

4.1 Hypotheses Explaining Gambling, and Insurance-Buying.¹

In recent papers, Friedman and Savage² have advanced a hypothesis, based on an assumed form of the utility function of money, attempting to explain why people may gamble, or buy insurance, or do both. Subsequently Markowitz³ advanced a modification of this theory, meeting certain difficulties inherent in the original theory. We shall discuss these theories in this section.

The central fact of a somewhat paradoxical nature in both gambling (at least in cases where there is a house 'cut'), and buying insurance is that the expected value of the money return in both these instances is negative. This appears paradoxical from the point of view of classical theories of gambling, which assumes that persons should take that action for which the expected value of the money return is the greatest. Clearly, the individual who gambles or buys insurance could choose a course which has a higher expectation of money return, by simply not gambling, or not gambling

1. It will become apparent in what follows that this theory applies only to insurance in which the buyer and the beneficiary are the same person, not, for example, to life insurance.

2. Friedman and Savage [9] and [10]

3. Markowitz [14]

the insurance. In modern theories of decision under risk, the person is supposed to choose that alternative which has the highest expected utility outcome, hence there is no contradiction with current utility theory in the fact that a person may not act to maximise expected money, and one may attempt, as Friedman and Savage and Markowitz do, to explain gambling and insurance buying by assuming that the utility function for money has a certain form.

Before proceeding to their theories, let us note that the fact that people do not play to maximise money, and that in some instances it seems utterly irrational to play this way, was noted in the 18th century in connection with the St. Petersburg paradox, which led Daniel Bernoulli to propose the first 'Bernoullian' utility scale. The St. Petersburg paradox concerns a game which is played in the following way. The 'house' allows the player to toss a fair coin as many times as necessary until it falls heads, then the house pays the player 2^n dollars. The question is, how much should the house charge the player to pay for the right to play this game? If the house is interested in making sure that its own expected money return is positive, then it should charge an amount slightly in excess of the expected value of the money to the player from playing the game. Conversely, if the player is interested in maximising the expected value of money return, he should be willing to pay any amount less than the expected value of the money return from the game for the privilege of playing it. However, it is easy to show that the expected value of money from this game is infinite, hence the player should be willing to pay any amount of money for the privilege of playing it. But to most people, even one thousand dollars would be too high a price to pay; the chance of even getting back the amount bet would be just one in two thousand. Several ingenious solutions were given to the paradox, most of them saving the principle that the

player should act to maximize his money expectation; however, Daniel Bernoulli's solution took the revolutionary tack of repudiating that principle, and proposed instead that players do indeed attempt to maximize a value, but that value is not proportional to money. Bernoulli solved the paradox by assuming that the value is proportional to its logarithm, from which it follows that that value of the game is exactly four dollars.

Even with Bernoulli's assumption, it is possible to modify the game in such a way that its expected value (or utility, in modern terms) is infinite: i.e. if the house pays not 2^n but 2^{2^n} dollars to the player if he tosses the coin n times before it falls heads, the value is then proportional to 2^n , and the expected value is infinite. In general, if the utility of money can be arbitrarily large, then it is possible to define a variant of the St. Petersburg game for which the expected value is infinite, and for which, therefore, the player should be willing to pay any amount to play. Since it seems unreasonable to be willing to pay an arbitrarily large amount to play any game, it can be argued that if the utility of money can be defined consistently at all, then it must be bounded above: i.e. if u is the utility function, and $u(x)$ is the utility of x dollars, there must be some number, say k , such that for all x , $u(x) < k$.

If the function u is plotted graphically, with x (the amount of money) on the horizontal axis, and $u(x)$ on the vertical axis, the above argument implies that there is a line above which the curve does not go (see Fig. 1).




Figure 1.

The theories of Friedman and Savage and Markowitz can be interpreted as giving other arguments like the above as to why the utility vs. money curve should have certain properties.

The first phenomenon which Friedman and Savage attempt to explain is gambling. They take as a typical case gambles in which there is a fairly small probability of winning a large amount, and a large probability of losing a small amount. Slot machines, roulette, and lotteries are among this type of gamble. All these games have the feature that the mathematical expectation of money winnings in playing them is negative, and is in fact measured by the 'house percentage.' Nevertheless, it is the case that people play them, and, even leaving aside the factor of excitement of participation (which we ruled out of consideration in our discussion of the axioms of utility), we may seek an explanation in terms of utility.

A typical gamble of the type referred to above may be represented in our formalism as follows. Let b be the amount the man bets, let w be the amount the man wins if he wins, let I be the amount of money he has at present, and let p be the probability of winning. Then $I + w$ is the total amount the man will have after playing if he wins, and $I - b$ is the total amount if he loses. He has probability p of ending up with $I + w$ and probability $1 - p$ of ending up with $I - b$: this is a risk outcome, and can be represented in our notation as: $\langle p(I+w), (1-p)(I-b) \rangle$.¹ The utility of this prospect is just $pu(I+w) + (1-p)u(I-b)$. If the man prefers to gamble, rather than not gamble and accept the certainty of remaining with the amount he has now, I , then it must be that

$$pu(I+w) + (1-p)u(I-b) > u(I)$$

On the other hand, we have postulated that the expected money gain from the gamble is negative. The expected money from gambling is just $p(I+w) + (1-p)(I-b)$ which we assume is less than I :

$$p(I+w) + (1-p)(I-b) < I$$

This situation is represented graphically in Figure 2. In this figure a straight line has been drawn between the points marking $u(I-b)$ and $u(I+w)$, and the expected utility point, $pu(I+w) + (1-p)u(I-b)$ is located on the line directly above the point on the x-axis marking the expected money value of the bet, $p(I+w) + (1-p)(I-b)$. The reader can easily convince himself that in general the expected utility of any probability combination of the extremes, $I+w$ and $I-b$, must lie on the straight line between the corresponding utility points, directly above the point on the x-axis indicating the expected money value of the probability combination.




Figure 2.

In Figure 2, the expected utility of the bet is shown to be greater than the utility of not betting, and the expected money value of the bet is shown to be less than the expected money value of not betting. A person who has a utility curve for which the three utilities $u(I-b)$, $u(I)$, and $u(I+w)$ have values as shown, could be expected to gamble in the situation described, and a utility curve of this type could be said to "explain" the phenomenon in question. If we look again at Figure 2, we note that a necessary condition that $u(I)$ be less than the expected utility of the gamble is that the point marked $u(I)$ in the figure lies below the line connecting the two points marked $u(I-b)$ and $u(I+w)$. This condition is met by a utility curve which is concave downward in the region under consideration, as shown in Figure 3.

Friedman and Savage

postulate a curve shaped as shown in Figure 3 as an explanation of gambling. In this figure I represents either current income, or customary income. The curve is shown concave downward above I , and this accounts for bets of the type considered, in which the possible gain is large, and the possible loss




Figure 3.

is small (note that in Figure 2, only three points on the utility curve are indicated, and the lowest one is only a short distance below I , whereas the highest one ($I+w$), is much higher than I).

The hypothesis that the utility curve is convex downward in some interval above I explains gambling, but does not prescribe what shape the curve is to have outside this region. Friedman and Savage propose that the utility curve is convex upward in the interval below I (current income) in order to explain the buying of insurance. Insurance buying is typified by paying a certain small amount, say r , for the security of having amount $I-p$ (which is the amount left from the present income after the insurance has been paid for). The insurance insures the man against a risk of losing a large amount, if an event with a small probability takes place, or else not losing anything in case the event doesn't take place. Let us suppose that the man loses d dollars if he is uninsured, and the event in question occurs, and that the event has probability p of occurring. Then the alternative of not taking the insurance has the risk outcome of getting $I-d$ dollars with probability p , and getting I dollars with probability $1-p$, and the alternative of buying the insurance has the certain outcome of getting $I-p$ dollars. The utility of the risk alternative is just $pu(I-d) + (1-p)u(I)$, and the utility of the second alternative is just $u(I-r)$. Since buying insurance is preferred to taking the chance, we must have:

$$pu(I-d) + (1-p)u(I) < u(I-r)$$

Furthermore, we have assumed that the expected value of the money to be gained from buying insurance is negative (at least, this seems to be the assumption that insurance companies operate under), hence:

$$p(I-d) + (1-p)I > I-r$$

In Figure 4 we represent this situation, and a curve is drawn which would explain the buying of insurance.




Figure 4.

It can be easily shown that in order for the two inequalities above to be satisfied, it is necessary that the point marked $u(I-r)$ lie above the line joining the points marked $u(I-d)$ and $u(I)$, and it is simplest to draw the utility curve as concave upward in the region in question.

In Figures 3 and 4, we have shown that utility curves which are concave downward to the right of I explain gambling, and curves which are concave upward to the left of I explain insurance buying: we can combine these into a single curve which explains both. However, before constructing the final curve, it is well to recall the discussion of the St. Petersburg paradox, in which it was argued that the utility curve must be bounded from

below. The curve resulting from all these arguments is shown in Figure 5. This curve seems to be the simplest type which is consistent with all of the facts discussed so far.

It is worthwhile to pause here and see if the curve thus drawn explains any other well-known facts other than the ones which it was originally constructed to explain.




Figure 5.

Friedman and Savage consider the factors influencing the distribution of prizes offered in lotteries. They note that almost all lotteries offer a graded series of prizes, starting with one or two very large prizes at the top, and working down to quite a few rather small prizes. They assume that the lottery operators attempt to construct the schedule of prizes in such a way that their profit from the lottery is a maximum subject to the restriction that the customers regard the tickets as worth the purchase price. This whole problem can be translated into utility terms in which lottery tickets represent risk outcomes with risk utilities which depend on the prizes offered and the probabilities of winning them, and the lottery operator seeks to adjust the prizes and probabilities in such a way that the utility of a ticket is greater than the utility of the purchase price, and at the same time the sum of the amount of the prizes is a minimum (and hence his profit is a maximum). Without going through the analysis here, we state

that the assumption that the utility curve is everywhere to the right of I convex downwards, instead of just in an initial interval, as shown in Figure 5, then the lottery ticket operator could make the most money by offering just a single very large prize, rather than by offering a number of prizes of varying amounts and varying probabilities. Therefore, the fact that lotteries do in fact offer a variety of prizes argues for the fact that the utility curve does not continue to bend upwards indefinitely to the right of I , and must instead start bending the other way again as it moves farther out.

Another fact cited by Markowitz is that people in general reject 'symmetrical' bets, that is, bets in which the amounts that can be lost or won are about the same (this is not supposed to extend to very small bets, in which it can be assumed that the amount of money involved is not important to the bettors). The fact that the curve as drawn in Figure 5 is symmetrical about the origin provides an explanation of this phenomenon. The reader can convince himself of this by representing the amounts to be won and lost at equal distances on either side of I , and connecting the corresponding utility points by a straight line, as was done in Figures 2 and 4. Bets which have a greater than 50% chance of losing will have utilities lying on this line to the left of its midpoint, hence below the x-axis, which represents the utility of I . Hence these bets will be rejected.

Friedman and Savage suggest that the utility curve may in fact be more complex than the one drawn in Figure 5: that it may instead have several 'humps', as shown in Figure 6. The curve of Figure 6 still explains all the facts mentioned so far, and there seems no reason to prefer one to the other. However, Friedman and Savage suggest that these 'steps' may in fact

represent discrete levels of aspiration for the individual, corresponding to definite social classes whose wealth corresponds to the different levels. At the top of each hump, there is a certain interval in which a large change in wealth carries little corresponding change of utility. Friedman and Savage say that this may be due to the fact that all the incomes in this




Figure 6.

interval are associated with one economic class, and that a change in wealth, as long as one remains in the same class, may not be important, whereas a change in wealth which carries a person from one class to another (corresponding to going from one step to another, over one of the steep intervals) may be regarded as much more important.

Before passing on to the next topic let us briefly note some possible objections to the theory just presented. First, as an explanation of gambling, it leaves out the very important factor of the excitement of participation. In the absence of any exact experimental data, it would seem that much of the type of gambling considered in this theory is of the kind in which the amount of money risked is quite small (at least for any one bet), and that the actual value of the money may be of comparable magnitude to the value of the excitement of the gamble. One is tempted to surmise that the purchasers of lottery tickets do not do so after sober consideration of the

relative values of the money bet and the prizes to be won, but act to a large extent on impulse. To argue that the amount of money spent on gambling may in total amount to a sizable portion of the gambler's income and hence that its value is large in comparison to the excitement of gambling is not to the point, since the stipulated interpretation of utility theory requires that it be applied to decisions made at a particular time. If it is assumed that utilities refer to average behavior, then the assumptions by which axiom A.4 (see page 49) was justified are violated, and it no longer follows that a Bernoullian utility function exists. Then, even if the total amount bet over a period of time is large, the amount bet at any given time by most people is small, and it seems likely that at the time the bet was made, one of the chief motivating factors was the thrill of betting; and our arguments for axiom A.7 imply that this must be a negligible factor if A.7 is to be satisfied, and a utility function exists.

In any event, the principal test which any theory must face is whether or not it succeeds in predicting a large variety of phenomena, and especially phenomena which it was not originally introduced to explain. Whether the above theory will meet this test we cannot say, but the criticisms suggest that if it is to be used with any precision, the basic interpretation will have to be more clearly defined. In the next section we discuss an experiment designed to test the theory, in discussing it we shall see one possible way of giving the basic concepts precise meanings.

4.2 The Mosteller-Nogee Experiment

In a recent paper, Mosteller and Nogee¹ have published the results of an experiment on gambling behavior which was intended as an empirical test of the Friedman-Savage theory discussed in Section 4.1.

1. Mosteller, Frederick, and Nogee, Philip [20]

This experiment consisted in running subjects through a series of gambles in which they were permitted either to bet 5¢, or not bet against various amounts of money offered at various odds by the experimenter. The game played was a variety of poker dice in which the experimenter rolled a "hand" of 5 dice and bet a certain sum, after which the subjects (each playing in turn) had the option of betting 5¢ and rolling the dice to try to beat the experimenter's hand, or not betting and passing the dice to the next subject.

According to the theory of Friedman and Savage, each subject should possess a "utility of money" curve, and should bet or not according as the expected utility of the bet offered by the experimenter is greater than or less than the utility of no change (i.e. not betting)¹. For the purposes of this experiment, the zero points of each person's utility scales were fixed at zero cents (i.e. at their state at the time of the bet), and the unit was chosen so that a loss of 5¢ had a utility of -1. With these two stipulations, each person's utility scale is fixed uniquely, and the utilities of every other gain or loss can be measured in terms of the utility of losing 5¢. According to the Friedman-Savage theory, once the zero point and unit of measurement have been chosen to determine the utility of any amount of money, say n cents, it is only necessary to find some probability p , such that the subject is indifferent between a bet which offers a probability p of winning n cents and $1-p$ of losing 5¢, and the alternative of not betting. If $u(n)$ is the utility of n cents, then the utility of a bet which offers a probability p of winning n cents, and $1-p$ of losing 5¢ is

$$pu(n) + (1-p)u(-5)$$

1. It appears that there is no operational meaning for the notion of "indifference" in this interpretation of utility theory. As we shall see, the experimental meanings at both "preference" and "indifference" as this experiment is actually carried out are considerably different from the interpretations given for these terms in Sect. 2.2.

and if this is held as indifferent to not betting, then

$$pu(n) + (1-p)u(-5) = u(0).$$

We have arbitrarily fixed the utility of 0¢ at 0 and utility of losing 5¢ as -1 so the above equation reduces to

$$pu(n) - (1-p) = 0,$$

or

$$u(n) = \frac{1-p}{p}$$

Morteller and Nogee's procedure was to assign definite operational meanings to the concepts of "preference" and "indifference" (which will be described below), then to determine some points on the curve of utility vs. money, using the formula given above, then to draw in a rough curve fitting the points plotted. Once a curve was drawn it was possible to test the Friedman-Savage theory, by presenting the subjects with various bets, somewhat more complicated than those which furnished the data from which the original curve was constructed, and noting whether their behavior in the new situations conformed to that predicted from the original curve. Thus, for a certain subject, they might plot several points on his utility vs. money curve using the above equation. Then draw in a rough curve of utility as shown in Figure 1. Once this is done, then




Figure 1.

if the Friedman-Savage hypothesis is correct, it should be possible to predict what the subject should do in all gambling situations in which the amounts of money involved fall within the range plotted in the figure.

It should be noted, of course, that the mere fact that a curve can be plotted using the formula:

$$u(n) = \frac{1-p}{p}$$

(where p is the probability at which the subject is indifferent between betting with a probability p of winning n cents and $1-p$ of losing $5¢$ and not betting) is not evidence tending to confirm the theory. Obviously there will be some probability for which the subject is indifferent in this situation, and putting that into the above formula, it is possible to calculate $u(n)$ in a mechanical way. The test of the theory is whether or not the subject chooses alternatives which maximize the expected value of the utilities thus calculated. Mosteller and Nogee tried two such tests, applying the information plotted in the original utility curve to try to predict behavior in new situations. The first test was to try to predict the behavior of subjects faced with "doublet" bets; that is, opportunities to make a single bet against two hands at the same time, where it is possible to win either one of two amounts of money, or both, or lose $5¢$. This is a different type of situation from that which provided the data on which the curve was based, but if the theory is correct, then the data contained in the plotted curve should predict the subject's behavior in the new situation. Hence the new situation furnishes a test for the theory. The doublet situation is represented formally as follows. Let p_1 and p_2 be the probabilities of beating the first and second hands respectively (assume that the first hand is higher than the second, hence the probability of beating it

is smaller: $p_1 < p_2$) and that n_1 and n_2 are the amounts to be won by beating hands 1 and 2 respectively. The probability of beating both the higher and lower hands and winning $n_1 + n_2$ cents is p_1 , the probability of beating the second hand but not the first hand and winning only n_2 cents is $p_2 - p_1$, and the probability of not beating either and losing 5¢ is $1 - p_2$. Hence, the utility of the doublet bet is:

$$p_1 u(n_1 + n_2) + (p_2 - p_1) u(n_2) + (1 - p_2) u(-5).$$

$u(n_1 + n_2)$, $u(n_2)$ and $u(-5)$ are all plotted on the utility curve, hence the utility of this bet can be calculated, and if the theory is correct, the subject should take the bet if this utility is greater than 0, be indifferent if the utility equals 0, and reject if the utility is less than 0.

A second test of the theory is afforded by "paired-comparison" situations. The principle idea is that the subject is forced to choose between one of two hands and money bets to bet against. Using the utility curve, the utility of each of the two bets offered by the experimenter can be calculated, and if the theory is correct, the subject should choose that bet with the highest utility. To describe this situation formally, suppose that the first bet offered by the experimenter is an amount n_1 on a hand which has probability p_1 of being beaten, and the second bet is n_2 on a hand which has probability p_2 of being beaten. If the subject bets against either hand, he must wager 5¢, hence the utility of the first bet is

$$p_1 u(n_1) + (1 - p_1) u(-5),$$

and the utility of the second is

$$p_2 u(n_2) + (1 - p_2) u(-5).$$

All these utilities are plotted on the curve already constructed so it is possible to calculate the utilities of these bets, and see whether the subject does in fact choose that with the highest utility.

The actual operational procedure for determining the points on the "original" utility curve (Figure 1) was as follows. A long series of trials was run during the course of which each subject had many opportunities of betting or not betting against each of the possible hands, and each of a number of offers on these hands made by the experimenter. Thus, one of the hands on which the experimenter made bets was four 4's and one 1, and among the many bets offered by the experimenter on that hand was 25¢, and during the course of the series this particular "hand", and the 25¢ bet by the experimenter were offered many times. At the end of the series, the proportion of times that a subject accepted a particular offer on a particular hand was calculated for each of the different offers on the hand, and was plotted as shown in Figure 2. Figure 2 shows the amounts offered on the hand on the horizontal axis, and the percentage of times that offer was accepted on the vertical axis. It was expected that for a fixed hand and subject, the higher the offer made, the greater the likelihood of acceptance,




Figure 2. Percentages of times bets of various amounts were accepted by subject X on hand A.

and that a curve drawn as in Fig. 2 would have approximately the "3" shape shown. This expectation was proved correct in all cases (except for one subject who left the experiment before its completion), although there was considerable variation in the steepness of the slopes of the steps of these curves. These curves were plotted for each subject and each hand, and the point at which they crossed the 50% level was taken to be the money offer on the hand for which the subject was indifferent.¹ For example, in the hypothetical curve drawn in Fig. 2, the indifference offer is approximately 37¢. If the probability of beating the hand is p , and the indifference offer is n_0 , then our formula allows us to calculate the utility at n i.e.,

$$u(n) = \frac{1-p}{p}.$$

Thus, each graph like that of Fig. 2 for a given subject determines one point on his utility curve, and by plotting these points it is possible to construct a curve like that of Fig. 1.²

Once the "basic" utility of money curves were plotted, it was possible to test the Friedman-Savage theory by applying the utilities of the basic curves to new situations. It is obvious that the theory cannot be expected to be completely successful in predicting the subject's choices because of the fact that the operational meaning given to "preference" is that the given alternative is chosen more than 50% of the time. But as long as it is possible for a subject to choose an alternative more than 50% of the time, but not all the time, then there must be instances in which he

1. Note that indifference is defined here as meaning that each of the alternatives is chosen 50% of the time; similarly, "preference" means that the preferred alternative is chosen more than 50% of the time. We shall discuss this interpretation in section 4.3.

2. The possibility that the subjects might not know the true probabilities of beating the various hands was ruled out by providing the subjects with lists giving the objective probabilities.

chooses an alternative which he does not prefer, and hence these are instances when the theory predicts he will choose one alternative (the preferred alternative), while he actually picks a different one. The fact that the subject's "S" curves, as illustrated in Fig. 2, have a non-vertical slope shows that counter-instances exist, in which the subject either chooses to bet, although the experimenter's offer is less than the indifference offer, or chooses not to bet, even though the experimenter's offer is greater than the indifference offer. Mosteller and Nogee make a comparative test of the Friedman-Savage theory by comparing the percentage of successful predictions from it with the percentage of successful predictions from a theory which assumes that the subjects act so as to maximize the expected value of money income. There they find that the Friedman-Savage theory is somewhat but not spectacularly more successful than the expected money hypothesis. Unfortunately, Mosteller and Nogee did not attempt to compare the Friedman-Savage theory with any other theories, and, as we shall see, there are reasons why the significance of their results is doubtful.

The fact that the Friedman-Savage theory turned out to be more successful than the expected money hypothesis should not seem surprising if it is recalled that both theories are very much alike in that they can both be interpreted as Bernoullian utility theories, and one (the Friedman-Savage) determines the utility of money empirically, whereas the other assumes that the utility of money curve is a straight line. It is natural that predictions based on a curve which is empirically determined should be more successful than predictions based on the a priori assumption that the utility curve is a straight line.

A second question which could be asked is whether the choice at the zero point of the utility scale as always being at the subject's present state, and of the loss of 5¢ as always being a change of one unit of utility is a correct interpretation. This amounts to assuming that what remains constant over time are the changes in utilities due to given changes of income. It would seem that it could be equally well argued that the utilities of various total amounts of money, or at least of total amounts of money on hand are what remain constant, and that the change in utility due to a given change in income (loss or gain of money) depends on what the total amounts are before and after the change. A second criticism, which Mosteller and Nogee comment on, is that it is possible that the subjects did not play each gamble separately, but might have played an over-all strategy with a view, not to maximizing their payoffs for each single gamble, but over a long series of gambles. It can be shown that the behavior which maximizes the expected payoff of a particular gamble is not necessarily the same as the behavior best calculated to maximize the total payoff due to a series of gambles of which the one in question is a part. Mosteller and Nogee note that, though the subjects seemed to be aware of long term strategic considerations¹, they did not follow their own convictions in these matters in actual play.

4.3 Problems of Interpretation and Confirmation: The Future of Bernoullian Utility Theory

The Mosteller-Nogee experiment brings up some problems inherent in any attempt at an empirical application of utility theory. We have alluded to these previously in section 1.3, but we are in a better position to discuss them now with a concrete example of an application before

1. As exemplified by the rule: "do not play a long shot when short of funds", which is justified by strategic considerations, but has no justification if the bet under consideration is considered in isolation.

us. The first thing to be noted is that Hosteller and Nogee interpret "preference" and "indifference" between two alternatives in terms of the relative frequencies with which one is chosen over the other. There is nothing wrong with this interpretation, except that the axioms for Bernoullian utilities were justified on a different basis, and there is good reason to believe that at least one of them - A.4 - should not hold under the relative frequency interpretation.

As the Hosteller-Nogee experiment shows, even aside from the problems involved with the relative frequency interpretation of preference, there is another fundamental difficulty involved in using utility theory predictively which does not arise under the "definition of rationality" interpretation. This is the difficulty of determining the individual's utility curves. If preferences between two alternatives is interpreted as meaning that one is always preferred to the other, then the individual's preference relation can be gradually constructed by observing him in a variety of choice situations. Even there, however, it is not possible to construct the utility curve precisely from a finite number of observations, and, as a matter of fact, it is not even possible to locate any points on it precisely, unless the subject has been observed in situations in which he is indifferent between certain alternatives. As the reader will recall, in order to locate some points on the subject's utility curves in the Hosteller-Nogee experiment, it was necessary to use an approximation to find values of money, n , for which the alternative of receiving n cents with probability p and losing 5¢ with probability $1-p$ is held as indifferent to not betting. Thus, at the present stage of the theory, even assuming that Bernoullian utility theory is "correct" in any of its empirical interpretations, its predictive usefulness is very much limited by

the difficulty of determining the utility functions of the people to whom it is to be applied.

Taken as a descriptive theory, utility theory, like decision theories in general, is a psychological theory. It is immediately evident that it cannot be precisely correct, because of the fact that the assumptions embodied in the axioms cannot be precisely satisfied. As a descriptive theory, it suffers from the further defect that in order for it to be applied, or tested, it requires the empirical determination of the utility function, which does not depend on only a small finite number of parameters as do many other theories. It would seem that unless some general psychological laws are discovered, relating to an individual's utility curves, utility theory will not be useful in predicting choice behavior, even though it may be approximately correct. Thus, the future of Bernoullian utility as a descriptive theory would appear to depend on two things: (1) whether or not it is nearly enough correct to make it workable to use in predictive situations; and (2), whether other psychological theories can be found from which utility curves can be inferred without the necessity of subjecting individuals to tests such as were used to determine the utility curves in the Mosteller-Nogee experiment.

BIBLIOGRAPHY

1. Allais, Maurice, "Le Comportement de l'Homme Rationnel Devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine." Econometrica, 21, (1953) pp. 503-546.
2. Arrow, K. J., "Alternative Approaches to the Theory of Choice in Risk-Taking Situations." Econometrica, 19, (1951) pp. 404-437.
3. ———, "Mathematical Models in Social Sciences," The Policy Sciences, Ed. Lasswell and Lerner, 1951.
4. Baumol, W. J., "The Neumann-Morgenstern Index: An Ordinalist View," Journal of Political Economy, February 1951.
5. Bishop, Robert L., "Consumer's Surplus and Cardinal Utility," Quarterly Journal of Economics, May 1943.
6. Blackwell, D. and Gershick, M. A., Theory of Games and Statistical Decisions, John Wiley and Sons, New York, 1954.
7. Davidson, D., McKinsey, J. C. G., and Suppes, P., Outlines of a Formal Theory of Value, I. Stanford Value Theory Project Report No-1, 1954.
8. Fisher, Irving, "A Statistical Method for Measuring Marginal Utility," Essays in Honor of T. B. Clark. Macmillan, 1929.
9. Friedman, M. and Savage, J. L., "The Expected Utility Hypothesis and the Measurability of Utility," Journal of Political Economy, December 1952, pp. 463-474.
10. ———, "The Utility Analysis of Choices Involving Risk," Journal of Political Economy, August 1948.
11. ———, "Methodes Nouvelles de Mesurer de l'Utilite Marginelle," Econometrica, January 1933.
12. Mallinraud, E., "Note on Von Neumann and Morgenstern's Strong Independence Axiom." Econometrica, October 1952.
13. Manns, Alan S., "The Strong Independence Assumption — Gasoline, Glenda and Probability Mixtures," Econometrica, October 1952.
14. Markowitz, Harry, "The Utility of Wealth," Journal of Political Economy, April 1952.
15. Marschok, Jacob, "Rational Behavior, Uncertain Prospects, and Measurable Utility," Econometrica, April 1950, Volume 18, pp. 111-141.

16. Milnor, T. W. and Herstein, I. M., "An Axiomatic Approach to Measurable Utility," Econometrica, April 1953.
17. Nash, John, "The Bargaining Problem," Econometrica, April 1950.
18. ———, "Two Person Cooperative Games," Econometrica, Vol. 21, #1, January 1953.
19. Neumann, John von and Morgenstern, Oskar, Theory of Games and Economic Behavior, 2nd Ed. New Jersey, 1947, Princeton University Press.
20. Negoe, P. and Monteller, P., "An Experimental Measure of Utility," Journal of Political Economy, October 1951.
21. Preston, M. G. and Baratta, P., "An Experimental Study of the Auction-Value of an Uncertain Outcome," American Journal of Psychology, 56, pp. 183-193, 1948.
22. Raiffa, Howard, "Arbitration Schemes for Generalized Two-Person Games," Contributions to the Theory of Games, Annual Mathematical Study 28, pp. 361-389, (1953).
23. Ramsey, Frank P., The Foundations of Mathematics and Other Logical Essays, Humanities Press, New York, 1950.
24. Robertson, D. H., "Utility and All That," Manchester School, May 1951.
25. Samuelson, Paul A., Foundations of Economic Analysis, Harvard University Press, Cambridge, 1953.
26. ———, "The Numerical Representation of Ordered Classification and the Concept of Utility," Rev. Economic Studies, October 1953.
27. ———, "Probability, Utility, and the Independence Axiom," Econometrica, October 1952.
28. Savage, L. T., "Note on Wold's: Ordinal Preferences or Cardinal Utility?" Econometrica, October 1952.
29. Stigler, G. J., "Development of Utility Theory," Journal of Political Economy, August and October 1950.
30. Vickrey, William, "Measuring Marginal Utility by Reactions to Risk," Econometrica, October 1945, Volume 13, #4.
31. Ville, T., "Sur les Conditions d'Existence d'une Ophelime Totale et d'un Indice Niveau des Prix," Annales de l'Universite de Lyon, Sec. A. Vol. 3, #9, 1946, pp. 32-39.
32. Viner, T., "Bentham and T. S. Mill," American Economic Review, XXXIX, 1949, pp. 860-82.

53. Viner, T., "The Utility Theory and its Critics," Journal of Political Economy, 1927.
54. Nold, H., "Ordinal Preferences or Cardinal Utility?" Econometrica, October 1952.

Unpublished Monographs

56. May, Kenneth O., "Transitivity, Utility and Aggregation of Preference Patterns," Unpublished, mimeographed paper.
57. Luce, R. D. and Raiffa, H., Survey of the Theory of Games. 7.R. No. 5, Bureau of Applied Social Research, May 1954.

Armed Services Technical Information Agency

Because of our limited supply, you are requested to return this copy WHEN IT HAS SERVED YOUR PURPOSE so that it may be made available to other requesters. Your cooperation will be appreciated.

AD

40840

NOTICE: WHEN GOVERNMENT OR OTHER DRAWINGS, SPECIFICATIONS OR OTHER DATA ARE USED FOR ANY PURPOSE OTHER THAN IN CONNECTION WITH A DEFINITELY RELATED GOVERNMENT PROCUREMENT OPERATION, THE U. S. GOVERNMENT THEREBY INCURS NO RESPONSIBILITY, NOR ANY OBLIGATION WHATSOEVER; AND THE FACT THAT THE GOVERNMENT MAY HAVE FORMULATED, FURNISHED, OR IN ANY WAY SUPPLIED THE SAID DRAWINGS, SPECIFICATIONS, OR OTHER DATA IS NOT TO BE REGARDED BY IMPLICATION OR OTHERWISE AS IN ANY MANNER LICENSING THE HOLDER OR ANY OTHER PERSON OR CORPORATION, OR CONVEYING ANY RIGHTS OR PERMISSION TO MANUFACTURE, USE OR SELL ANY PATENTED INVENTION THAT MAY IN ANY WAY BE RELATED THERETO.

Reproduced by

DOCUMENT SERVICE CENTER

KNOTT BUILDING, DAYTON, 2, OHIO

UNCLASSIFIED